



API Reference

Amazon S3



Amazon S3: API Reference

Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

Welcome	1
Amazon S3	1
Amazon S3 Control	1
Amazon S3 Files	2
Amazon S3 on Outposts	2
Amazon S3 Tables	2
Amazon S3 Vectors	2
Actions	3
Amazon S3	13
AbortMultipartUpload	18
CompleteMultipartUpload	25
CopyObject	44
CreateBucket	79
CreateBucketMetadataConfiguration	95
CreateBucketMetadataTableConfiguration	102
CreateMultipartUpload	106
CreateSession	133
DeleteBucket	141
DeleteBucketAnalyticsConfiguration	145
DeleteBucketCors	148
DeleteBucketEncryption	151
DeleteBucketIntelligentTieringConfiguration	155
DeleteBucketInventoryConfiguration	158
DeleteBucketLifecycle	162
DeleteBucketMetadataConfiguration	166
DeleteBucketMetadataTableConfiguration	169
DeleteBucketMetricsConfiguration	172
DeleteBucketOwnershipControls	176
DeleteBucketPolicy	179
DeleteBucketReplication	183
DeleteBucketTagging	186
DeleteBucketWebsite	189
DeleteObject	192
DeleteObjectAnnotation	203

DeleteObjects	207
DeleteObjectTagging	226
DeletePublicAccessBlock	230
GetBucketAbac	233
GetBucketAccelerateConfiguration	235
GetBucketAcl	240
GetBucketAnalyticsConfiguration	245
GetBucketCors	250
GetBucketEncryption	255
GetBucketIntelligentTieringConfiguration	260
GetBucketInventoryConfiguration	264
GetBucketLifecycle	271
GetBucketLifecycleConfiguration	276
GetBucketLocation	284
GetBucketLogging	289
GetBucketMetadataConfiguration	294
GetBucketMetadataTableConfiguration	298
GetBucketMetricsConfiguration	302
GetBucketNotification	310
GetBucketNotificationConfiguration	314
GetBucketOwnershipControls	320
GetBucketPolicy	325
GetBucketPolicyStatus	330
GetBucketReplication	334
GetBucketRequestPayment	340
GetBucketTagging	344
GetBucketVersioning	348
GetBucketWebsite	352
GetObject	357
GetObjectAcl	385
GetObjectAnnotation	392
GetObjectAttributes	398
GetObjectLegalHold	410
GetObjectLockConfiguration	414
GetObjectRetention	418
GetObjectTagging	422

GetObjectTorrent	427
GetPublicAccessBlock	431
HeadBucket	436
HeadObject	444
ListBucketAnalyticsConfigurations	470
ListBucketIntelligentTieringConfigurations	477
ListBucketInventoryConfigurations	481
ListBucketMetricsConfigurations	489
ListBuckets	496
ListDirectoryBuckets	504
ListMultipartUploads	508
ListObjectAnnotations	525
ListObjects	531
ListObjectsV2	544
ListObjectVersions	563
ListParts	581
PutBucketAbac	594
PutBucketAccelerateConfiguration	597
PutBucketAcl	602
PutBucketAnalyticsConfiguration	613
PutBucketCors	619
PutBucketEncryption	625
PutBucketIntelligentTieringConfiguration	633
PutBucketInventoryConfiguration	638
PutBucketLifecycle	647
PutBucketLifecycleConfiguration	656
PutBucketLogging	670
PutBucketMetricsConfiguration	677
PutBucketNotification	685
PutBucketNotificationConfiguration	689
PutBucketOwnershipControls	698
PutBucketPolicy	704
PutBucketReplication	711
PutBucketRequestPayment	720
PutBucketTagging	724
PutBucketVersioning	729

PutBucketWebsite	736
PutObject	746
PutObjectAcl	781
PutObjectAnnotation	794
PutObjectLegalHold	802
PutObjectLockConfiguration	807
PutObjectRetention	812
PutObjectTagging	817
PutPublicAccessBlock	823
RenameObject	830
RestoreObject	837
SelectObjectContent	849
UpdateBucketMetadataAnnotationTableConfiguration	861
UpdateBucketMetadataInventoryTableConfiguration	867
UpdateBucketMetadataJournalTableConfiguration	871
UpdateObjectEncryption	874
UploadPart	882
UploadPartCopy	898
WriteGetObjectResponse	917
Amazon S3 Control	928
AssociateAccessGrantsIdentityCenter	932
CreateAccessGrant	935
CreateAccessGrantsInstance	942
CreateAccessGrantsLocation	947
CreateAccessPoint	952
CreateAccessPointForObjectLambda	959
CreateBucket	963
CreateJob	970
CreateMultiRegionAccessPoint	981
CreateStorageLensGroup	985
DeleteAccessGrant	990
DeleteAccessGrantsInstance	992
DeleteAccessGrantsInstanceResourcePolicy	994
DeleteAccessGrantsLocation	996
DeleteAccessPoint	998
DeleteAccessPointForObjectLambda	1001

DeleteAccessPointPolicy	1003
DeleteAccessPointPolicyForObjectLambda	1006
DeleteAccessPointScope	1008
DeleteBucket	1010
DeleteBucketLifecycleConfiguration	1013
DeleteBucketPolicy	1016
DeleteBucketReplication	1020
DeleteBucketTagging	1024
DeleteJobTagging	1027
DeleteMultiRegionAccessPoint	1030
DeletePublicAccessBlock	1034
DeleteStorageLensConfiguration	1036
DeleteStorageLensConfigurationTagging	1039
DeleteStorageLensGroup	1042
DescribeJob	1044
DescribeMultiRegionAccessPointOperation	1053
DissociateAccessGrantsIdentityCenter	1057
GetAccessGrant	1059
GetAccessGrantsInstance	1064
GetAccessGrantsInstanceForPrefix	1068
GetAccessGrantsInstanceResourcePolicy	1071
GetAccessGrantsLocation	1074
GetAccessPoint	1078
GetAccessPointConfigurationForObjectLambda	1084
GetAccessPointForObjectLambda	1087
GetAccessPointPolicy	1090
GetAccessPointPolicyForObjectLambda	1093
GetAccessPointPolicyStatus	1097
GetAccessPointPolicyStatusForObjectLambda	1100
GetAccessPointScope	1103
GetBucket	1106
GetBucketLifecycleConfiguration	1110
GetBucketPolicy	1115
GetBucketReplication	1119
GetBucketTagging	1125
GetBucketVersioning	1129

GetDataAccess	1134
GetJobTagging	1139
GetMultiRegionAccessPoint	1142
GetMultiRegionAccessPointPolicy	1145
GetMultiRegionAccessPointPolicyStatus	1148
GetMultiRegionAccessPointRoutes	1151
GetPublicAccessBlock	1154
GetStorageLensConfiguration	1158
GetStorageLensConfigurationTagging	1166
GetStorageLensGroup	1169
ListAccessGrants	1174
ListAccessGrantsInstances	1178
ListAccessGrantsLocations	1181
ListAccessPoints	1184
ListAccessPointsForDirectoryBuckets	1189
ListAccessPointsForObjectLambda	1193
ListCallerAccessGrants	1197
ListJobs	1201
ListMultiRegionAccessPoints	1205
ListRegionalBuckets	1209
ListStorageLensConfigurations	1213
ListStorageLensGroups	1216
ListTagsForResource	1219
PutAccessGrantsInstanceResourcePolicy	1223
PutAccessPointConfigurationForObjectLambda	1227
PutAccessPointPolicy	1230
PutAccessPointPolicyForObjectLambda	1234
PutAccessPointScope	1238
PutBucketLifecycleConfiguration	1241
PutBucketPolicy	1246
PutBucketReplication	1251
PutBucketTagging	1259
PutBucketVersioning	1264
PutJobTagging	1269
PutMultiRegionAccessPointPolicy	1273
PutPublicAccessBlock	1277

PutStorageLensConfiguration	1281
PutStorageLensConfigurationTagging	1287
SubmitMultiRegionAccessPointRoutes	1290
TagResource	1295
UntagResource	1299
UpdateAccessGrantsLocation	1302
UpdateJobPriority	1307
UpdateJobStatus	1311
UpdateStorageLensGroup	1315
Amazon S3 Files	1318
CreateAccessPoint	1320
CreateFileSystem	1328
CreateMountTarget	1337
DeleteAccessPoint	1345
DeleteFileSystem	1348
DeleteFileSystemPolicy	1351
DeleteMountTarget	1354
GetAccessPoint	1357
GetFileSystem	1363
GetFileSystemPolicy	1369
GetMountTarget	1372
GetSynchronizationConfiguration	1377
ListAccessPoints	1381
ListFileSystems	1385
ListMountTargets	1388
ListTagsForResource	1392
PutFileSystemPolicy	1396
PutSynchronizationConfiguration	1399
TagResource	1403
UntagResource	1406
UpdateMountTarget	1409
Amazon S3 on Outposts	1414
CreateEndpoint	1415
DeleteEndpoint	1420
ListEndpoints	1423
ListOutpostsWithS3	1427

ListSharedEndpoints	1430
Amazon S3 Tables	1433
CreateNamespace	1436
CreateTable	1440
CreateTableBucket	1446
DeleteNamespace	1451
DeleteTable	1454
DeleteTableBucket	1457
DeleteTableBucketEncryption	1460
DeleteTableBucketMetricsConfiguration	1463
DeleteTableBucketPolicy	1466
DeleteTableBucketReplication	1469
DeleteTablePolicy	1472
DeleteTableReplication	1475
GetNamespace	1478
GetTable	1482
GetTableBucket	1489
GetTableBucketEncryption	1493
GetTableBucketMaintenanceConfiguration	1496
GetTableBucketMetricsConfiguration	1500
GetTableBucketPolicy	1503
GetTableBucketReplication	1506
GetTableBucketStorageClass	1510
GetTableEncryption	1513
GetTableMaintenanceConfiguration	1517
GetTableMaintenanceJobStatus	1521
GetTableMetadataLocation	1525
GetTablePolicy	1529
GetTableRecordExpirationConfiguration	1532
GetTableRecordExpirationJobStatus	1535
GetTableReplication	1539
GetTableReplicationStatus	1543
GetTableStorageClass	1547
ListNamespaces	1551
ListTableBuckets	1555
ListTables	1559

ListTagsForResource	1563
PutTableBucketEncryption	1566
PutTableBucketMaintenanceConfiguration	1569
PutTableBucketMetricsConfiguration	1572
PutTableBucketPolicy	1575
PutTableBucketReplication	1578
PutTableBucketStorageClass	1582
PutTableMaintenanceConfiguration	1585
PutTablePolicy	1589
PutTableRecordExpirationConfiguration	1593
PutTableReplication	1596
RenameTable	1600
TagResource	1604
UntagResource	1607
UpdateTableMetadataLocation	1610
Amazon S3 Vectors	1614
CreateIndex	1616
CreateVectorBucket	1622
DeleteIndex	1627
DeleteVectorBucket	1631
DeleteVectorBucketPolicy	1635
DeleteVectors	1638
GetIndex	1643
GetVectorBucket	1647
GetVectorBucketPolicy	1651
GetVectors	1655
ListIndexes	1661
ListTagsForResource	1666
ListVectorBuckets	1670
ListVectors	1674
PutVectorBucketPolicy	1681
PutVectors	1685
QueryVectors	1690
TagResource	1697
UntagResource	1701
Data Types	1705

Amazon S3	1718
AbacStatus	1724
AbortIncompleteMultipartUpload	1725
AccelerateConfiguration	1726
AccessControlPolicy	1727
AccessControlTranslation	1728
AnalyticsAndOperator	1729
AnalyticsConfiguration	1730
AnalyticsExportDestination	1732
AnalyticsFilter	1733
AnalyticsS3BucketDestination	1735
AnnotationEntry	1737
AnnotationTableConfiguration	1739
AnnotationTableConfigurationResult	1741
AnnotationTableConfigurationUpdates	1743
BlockedEncryptionTypes	1745
Bucket	1747
BucketInfo	1749
BucketLifecycleConfiguration	1751
BucketLoggingStatus	1752
Checksum	1753
CloudFunctionConfiguration	1757
CommonPrefix	1760
CompletedMultipartUpload	1761
CompletedPart	1762
Condition	1766
ContinuationEvent	1768
CopyObjectResult	1769
CopyPartResult	1773
CORSConfiguration	1777
CORSRule	1778
CreateBucketConfiguration	1780
CSVInput	1783
CSVOutput	1786
DefaultRetention	1788
Delete	1790

DeletedObject	1792
DeleteMarkerEntry	1794
DeleteMarkerReplication	1796
Destination	1798
DestinationResult	1801
Encryption	1803
EncryptionConfiguration	1805
EndEvent	1806
Error	1807
ErrorDetails	1822
ErrorDocument	1827
EventBridgeConfiguration	1828
EventHoldDuration	1829
ExistingObjectReplication	1830
FilterRule	1831
GetBucketMetadataConfigurationResult	1833
GetBucketMetadataTableConfigurationResult	1834
GetObjectAttributesParts	1836
GlacierJobParameters	1838
Grant	1839
Grantee	1840
IndexDocument	1842
Initiator	1843
InputSerialization	1845
IntelligentTieringAndOperator	1847
IntelligentTieringConfiguration	1848
IntelligentTieringFilter	1850
InventoryConfiguration	1852
InventoryDestination	1855
InventoryEncryption	1856
InventoryFilter	1857
InventoryS3BucketDestination	1858
InventorySchedule	1860
InventoryTableConfiguration	1861
InventoryTableConfigurationResult	1862
InventoryTableConfigurationUpdates	1864

JournalTableConfiguration	1865
JournalTableConfigurationResult	1866
JournalTableConfigurationUpdates	1868
JSONInput	1869
JSONOutput	1870
LambdaFunctionConfiguration	1871
LifecycleConfiguration	1873
LifecycleExpiration	1874
LifecycleRule	1876
LifecycleRuleAndOperator	1880
LifecycleRuleFilter	1882
LocationInfo	1884
LoggingEnabled	1886
MetadataConfiguration	1888
MetadataConfigurationResult	1889
MetadataEntry	1891
MetadataTableConfiguration	1892
MetadataTableConfigurationResult	1893
MetadataTableEncryptionConfiguration	1894
Metrics	1895
MetricsAndOperator	1896
MetricsConfiguration	1898
MetricsFilter	1900
MultipartUpload	1902
NoncurrentVersionExpiration	1905
NoncurrentVersionTransition	1907
NotificationConfiguration	1909
NotificationConfigurationDeprecated	1911
NotificationConfigurationFilter	1913
Object	1914
ObjectEncryption	1918
ObjectIdentifier	1919
ObjectLockConfiguration	1922
ObjectLockLegalHold	1923
ObjectLockRetention	1924
ObjectLockRule	1926

ObjectPart	1927
ObjectVersion	1931
OutputLocation	1934
OutputSerialization	1935
Owner	1936
OwnershipControls	1937
OwnershipControlsRule	1938
ParquetInput	1940
Part	1941
PartitionedPrefix	1945
PolicyStatus	1946
Progress	1947
ProgressEvent	1948
PublicAccessBlockConfiguration	1949
QueueConfiguration	1951
QueueConfigurationDeprecated	1953
RecordExpiration	1956
RecordsEvent	1957
Redirect	1958
RedirectAllRequestsTo	1960
ReplicaModifications	1961
ReplicationConfiguration	1962
ReplicationRule	1963
ReplicationRuleAndOperator	1967
ReplicationRuleFilter	1968
ReplicationTime	1970
ReplicationTimeValue	1971
RequestPaymentConfiguration	1972
RequestProgress	1973
RestoreRequest	1974
RestoreStatus	1977
RoutingRule	1979
Rule	1980
S3KeyFilter	1983
S3Location	1984
S3TablesDestination	1986

S3TablesDestinationResult	1988
ScanRange	1990
SelectObjectContentEventStream	1991
SelectParameters	1993
ServerSideEncryptionByDefault	1995
ServerSideEncryptionConfiguration	1998
ServerSideEncryptionRule	1999
SessionCredentials	2002
SimplePrefix	2004
SourceSelectionCriteria	2005
SSEKMS	2007
SseKmsEncryptedObjects	2008
SSEKMSEncryption	2009
SSES3	2011
Stats	2012
StatsEvent	2013
StorageClassAnalysis	2014
StorageClassAnalysisDataExport	2015
Tag	2016
Tagging	2017
TargetGrant	2018
TargetObjectKeyFormat	2019
Tiering	2020
TopicConfiguration	2021
TopicConfigurationDeprecated	2023
Transition	2026
VersioningConfiguration	2028
WebsiteConfiguration	2029
Amazon S3 Control	2030
AbortIncompleteMultipartUpload	2036
AccessControlTranslation	2037
AccessGrantsLocationConfiguration	2038
AccessPoint	2040
AccountLevel	2043
ActivityMetrics	2045
AdvancedCostOptimizationMetrics	2046

AdvancedDataProtectionMetrics	2047
AdvancedPerformanceMetrics	2048
AsyncErrorDetails	2049
AsyncOperation	2051
AsyncRequestParameters	2053
AsyncResponseDetails	2054
AwsLambdaTransformation	2055
BucketLevel	2056
CloudWatchMetrics	2058
CreateBucketConfiguration	2059
CreateMultiRegionAccessPointInput	2061
Credentials	2063
DeleteMarkerReplication	2065
DeleteMultiRegionAccessPointInput	2066
Destination	2067
DetailedStatusCodesMetrics	2070
DSSEKMSFilter	2071
EncryptionConfiguration	2072
EstablishedMultiRegionAccessPointPolicy	2073
Exclude	2074
ExistingObjectReplication	2075
GeneratedManifestEncryption	2076
Grantee	2077
Include	2079
JobDescriptor	2080
JobFailure	2085
JobListDescriptor	2086
JobManifest	2089
JobManifestGenerator	2090
JobManifestGeneratorFilter	2091
JobManifestLocation	2094
JobManifestSpec	2096
JobOperation	2097
JobProgressSummary	2101
JobReport	2103
JobTimers	2105

KeyNameConstraint	2106
LambdaInvokeOperation	2108
LifecycleConfiguration	2110
LifecycleExpiration	2111
LifecycleRule	2113
LifecycleRuleAndOperator	2116
LifecycleRuleFilter	2118
ListAccessGrantEntry	2120
ListAccessGrantsInstanceEntry	2123
ListAccessGrantsLocationsEntry	2126
ListCallerAccessGrantsEntry	2128
ListStorageLensConfigurationEntry	2130
ListStorageLensGroupEntry	2132
MatchObjectAge	2134
MatchObjectSize	2135
Metrics	2136
MultiRegionAccessPointPolicyDocument	2137
MultiRegionAccessPointRegionalResponse	2138
MultiRegionAccessPointReport	2139
MultiRegionAccessPointRoute	2142
MultiRegionAccessPointsAsyncResponse	2144
NoncurrentVersionExpiration	2145
NoncurrentVersionTransition	2146
NotSSEFilter	2147
ObjectEncryption	2148
ObjectEncryptionFilter	2149
ObjectLambdaAccessPoint	2151
ObjectLambdaAccessPointAlias	2153
ObjectLambdaConfiguration	2155
ObjectLambdaContentTransformation	2157
ObjectLambdaTransformationConfiguration	2158
PolicyStatus	2159
PrefixLevel	2160
PrefixLevelStorageMetrics	2161
ProposedMultiRegionAccessPointPolicy	2162
PublicAccessBlockConfiguration	2163

PutMultiRegionAccessPointPolicyInput	2165
Region	2166
RegionalBucket	2167
RegionReport	2169
ReplicaModifications	2171
ReplicationConfiguration	2172
ReplicationRule	2173
ReplicationRuleAndOperator	2176
ReplicationRuleFilter	2177
ReplicationTime	2179
ReplicationTimeValue	2181
S3AccessControlList	2182
S3AccessControlPolicy	2183
S3BucketDestination	2184
S3ComputeObjectChecksumOperation	2186
S3CopyObjectOperation	2187
S3DeleteObjectTaggingOperation	2195
S3GeneratedManifestDescriptor	2196
S3Grant	2197
S3Grantee	2198
S3InitiateRestoreObjectOperation	2199
S3JobManifestGenerator	2201
S3ManifestOutputLocation	2203
S3ObjectLockEventHoldDuration	2205
S3ObjectLockLegalHold	2206
S3ObjectLockRetentionEventHoldDuration	2207
S3ObjectMetadata	2208
S3ObjectOwner	2211
S3ReplicateObjectOperation	2212
S3Retention	2213
S3SetObjectAclOperation	2215
S3SetObjectLegalHoldOperation	2216
S3SetObjectRetentionOperation	2217
S3SetObjectTaggingOperation	2219
S3Tag	2220
S3UpdateObjectEncryptionOperation	2221

S3UpdateObjectEncryptionSSEKMS	2222
Scope	2224
SelectionCriteria	2225
SourceSelectionCriteria	2227
SSECFilter	2229
SSEKMS	2230
SseKmsEncryptedObjects	2231
SSEKMSEncryption	2232
SSEKMSFilter	2233
SSES3	2234
SSES3Encryption	2235
SSES3Filter	2236
StorageLensAwsOrg	2237
StorageLensConfiguration	2238
StorageLensDataExport	2241
StorageLensDataExportEncryption	2243
StorageLensExpandedPrefixesDataExport	2244
StorageLensGroup	2245
StorageLensGroupAndOperator	2247
StorageLensGroupFilter	2249
StorageLensGroupLevel	2251
StorageLensGroupLevelSelectionCriteria	2252
StorageLensGroupOrOperator	2254
StorageLensTableDestination	2256
StorageLensTag	2257
Tag	2258
Tagging	2260
Transition	2261
VersioningConfiguration	2263
VpcConfiguration	2264
Amazon S3 Files	2264
CreationPermissions	2266
ExpirationDataRule	2268
ImportDataRule	2269
ListAccessPointsDescription	2271
ListFileSystemsDescription	2274

ListMountTargetsDescription	2277
PosixUser	2280
RootDirectory	2282
Tag	2283
Amazon S3 on Outposts	2284
Endpoint	2285
FailedReason	2288
NetworkInterface	2289
Outpost	2290
Amazon S3 Tables	2291
EncryptionConfiguration	2293
IcebergCompactionSettings	2295
IcebergMetadata	2296
IcebergPartitionField	2298
IcebergPartitionSpec	2300
IcebergSchema	2301
IcebergSchemaV2	2302
IcebergSnapshotManagementSettings	2304
IcebergSortField	2305
IcebergSortOrder	2307
IcebergUnreferencedFileRemovalSettings	2308
LastSuccessfulReplicatedUpdate	2309
ManagedTableInformation	2310
NamespaceSummary	2311
ReplicationDestination	2313
ReplicationDestinationStatusModel	2314
ReplicationInformation	2316
SchemaField	2317
SchemaV2Field	2319
StorageClassConfiguration	2321
TableBucketMaintenanceConfigurationValue	2322
TableBucketMaintenanceSettings	2323
TableBucketReplicationConfiguration	2324
TableBucketReplicationRule	2325
TableBucketSummary	2326
TableMaintenanceConfigurationValue	2328

TableMaintenanceJobStatusValue	2329
TableMaintenanceSettings	2330
TableMetadata	2331
TableRecordExpirationConfigurationValue	2332
TableRecordExpirationJobMetrics	2333
TableRecordExpirationSettings	2334
TableReplicationConfiguration	2335
TableReplicationRule	2336
TableSummary	2337
Amazon S3 Vectors	2339
EncryptionConfiguration	2340
GetOutputVector	2342
Index	2344
IndexSummary	2347
ListOutputVector	2349
MetadataConfiguration	2351
PutInputVector	2352
QueryOutputVector	2354
ValidationExceptionField	2355
VectorBucket	2356
VectorBucketSummary	2358
VectorData	2360

Welcome

Amazon S3

Note

For information about using the Amazon S3 API—including authentication, signing requests, code examples, and error handling—see the [Amazon S3 Developer Guide](#).

Welcome to the *Amazon S3 API Reference*. This guide explains the Amazon Simple Storage Service (Amazon S3) application programming interface (API).

Welcome to the *Amazon S3 API Reference*. This guide explains the Amazon Simple Storage Service (Amazon S3) application programming interface (API).

You can use any toolkit that supports HTTP to use the REST API. You can even use a browser to fetch objects, as long as they are anonymously readable.

The REST API uses the standard HTTP headers and status codes, so that standard browsers and toolkits work as expected. In some areas, we have added functionality to HTTP (for example, we added headers to support access control). In these cases, we have done our best to add the new functionality in a way that matched the style of standard HTTP usage.

The current version of the Amazon S3 API is 2006-03-01.

Amazon S3 supports the REST API.

Note

Support for SOAP over HTTP is deprecated, but it is still available over HTTPS. However, new Amazon S3 features will not be supported for SOAP. We recommend that you use either this REST API or the AWS SDKs.

Amazon S3 Control

AWS S3 Control provides access to Amazon S3 control plane actions.

Amazon S3 Files

S3 Files makes S3 buckets accessible as high-performance file systems powered by EFS. This service enables file system interface access to S3 data with sub-millisecond latencies through mount targets, supporting AI/ML workloads, media processing, and hybrid storage workflows that require both file system and object storage access to the same data.

Amazon S3 on Outposts

Amazon S3 on Outposts provides access to S3 on Outposts operations.

Amazon S3 Tables

An Amazon S3 table represents a structured dataset consisting of tabular data in [Apache Parquet](#) format and related metadata. This data is stored inside an S3 table as a subresource. All tables in a table bucket are stored in the [Apache Iceberg](#) table format. Through integration with the [AWS Glue Data Catalog](#) you can interact with your tables using AWS analytics services, such as [Amazon Athena](#) and [Amazon Redshift](#). Amazon S3 manages maintenance of your tables through automatic file compaction and snapshot management. For more information, see [Amazon S3 table buckets](#).

Amazon S3 Vectors

Amazon S3 vector buckets are a bucket type to store and search vectors with sub-second search times. They are designed to provide dedicated API operations for you to interact with vectors to do similarity search. Within a vector bucket, you use a vector index to organize and logically group your vector data. When you make a write or read request, you direct it to a single vector index. You store your vector data as vectors. A vector contains a key (a name that you assign), a multi-dimensional vector, and, optionally, metadata that describes a vector. The key uniquely identifies the vector in a vector index.

Actions

The following actions are supported by Amazon S3:

- [AbortMultipartUpload](#)
- [CompleteMultipartUpload](#)
- [CopyObject](#)
- [CreateBucket](#)
- [CreateBucketMetadataConfiguration](#)
- [CreateBucketMetadataTableConfiguration](#)
- [CreateMultipartUpload](#)
- [CreateSession](#)
- [DeleteBucket](#)
- [DeleteBucketAnalyticsConfiguration](#)
- [DeleteBucketCors](#)
- [DeleteBucketEncryption](#)
- [DeleteBucketIntelligentTieringConfiguration](#)
- [DeleteBucketInventoryConfiguration](#)
- [DeleteBucketLifecycle](#)
- [DeleteBucketMetadataConfiguration](#)
- [DeleteBucketMetadataTableConfiguration](#)
- [DeleteBucketMetricsConfiguration](#)
- [DeleteBucketOwnershipControls](#)
- [DeleteBucketPolicy](#)
- [DeleteBucketReplication](#)
- [DeleteBucketTagging](#)
- [DeleteBucketWebsite](#)
- [DeleteObject](#)
- [DeleteObjectAnnotation](#)
- [DeleteObjects](#)
- [DeleteObjectTagging](#)

- [DeletePublicAccessBlock](#)
- [GetBucketAbac](#)
- [GetBucketAccelerateConfiguration](#)
- [GetBucketAcl](#)
- [GetBucketAnalyticsConfiguration](#)
- [GetBucketCors](#)
- [GetBucketEncryption](#)
- [GetBucketIntelligentTieringConfiguration](#)
- [GetBucketInventoryConfiguration](#)
- [GetBucketLifecycle](#)
- [GetBucketLifecycleConfiguration](#)
- [GetBucketLocation](#)
- [GetBucketLogging](#)
- [GetBucketMetadataConfiguration](#)
- [GetBucketMetadataTableConfiguration](#)
- [GetBucketMetricsConfiguration](#)
- [GetBucketNotification](#)
- [GetBucketNotificationConfiguration](#)
- [GetBucketOwnershipControls](#)
- [GetBucketPolicy](#)
- [GetBucketPolicyStatus](#)
- [GetBucketReplication](#)
- [GetBucketRequestPayment](#)
- [GetBucketTagging](#)
- [GetBucketVersioning](#)
- [GetBucketWebsite](#)
- [GetObject](#)
- [GetObjectAcl](#)
- [GetObjectAnnotation](#)
- [GetObjectAttributes](#)

- [GetObjectLegalHold](#)
- [GetObjectLockConfiguration](#)
- [GetObjectRetention](#)
- [GetObjectTagging](#)
- [GetObjectTorrent](#)
- [GetPublicAccessBlock](#)
- [HeadBucket](#)
- [HeadObject](#)
- [ListBucketAnalyticsConfigurations](#)
- [ListBucketIntelligentTieringConfigurations](#)
- [ListBucketInventoryConfigurations](#)
- [ListBucketMetricsConfigurations](#)
- [ListBuckets](#)
- [ListDirectoryBuckets](#)
- [ListMultipartUploads](#)
- [ListObjectAnnotations](#)
- [ListObjects](#)
- [ListObjectsV2](#)
- [ListObjectVersions](#)
- [ListParts](#)
- [PutBucketAbac](#)
- [PutBucketAccelerateConfiguration](#)
- [PutBucketAcl](#)
- [PutBucketAnalyticsConfiguration](#)
- [PutBucketCors](#)
- [PutBucketEncryption](#)
- [PutBucketIntelligentTieringConfiguration](#)
- [PutBucketInventoryConfiguration](#)
- [PutBucketLifecycle](#)
- [PutBucketLifecycleConfiguration](#)

- [PutBucketLogging](#)
- [PutBucketMetricsConfiguration](#)
- [PutBucketNotification](#)
- [PutBucketNotificationConfiguration](#)
- [PutBucketOwnershipControls](#)
- [PutBucketPolicy](#)
- [PutBucketReplication](#)
- [PutBucketRequestPayment](#)
- [PutBucketTagging](#)
- [PutBucketVersioning](#)
- [PutBucketWebsite](#)
- [PutObject](#)
- [PutObjectAcl](#)
- [PutObjectAnnotation](#)
- [PutObjectLegalHold](#)
- [PutObjectLockConfiguration](#)
- [PutObjectRetention](#)
- [PutObjectTagging](#)
- [PutPublicAccessBlock](#)
- [RenameObject](#)
- [RestoreObject](#)
- [SelectObjectContent](#)
- [UpdateBucketMetadataAnnotationTableConfiguration](#)
- [UpdateBucketMetadataInventoryTableConfiguration](#)
- [UpdateBucketMetadataJournalTableConfiguration](#)
- [UpdateObjectEncryption](#)
- [UploadPart](#)
- [UploadPartCopy](#)
- [WriteGetObjectResponse](#)

The following actions are supported by Amazon S3 Control:

- [AssociateAccessGrantsIdentityCenter](#)
- [CreateAccessGrant](#)
- [CreateAccessGrantsInstance](#)
- [CreateAccessGrantsLocation](#)
- [CreateAccessPoint](#)
- [CreateAccessPointForObjectLambda](#)
- [CreateBucket](#)
- [CreateJob](#)
- [CreateMultiRegionAccessPoint](#)
- [CreateStorageLensGroup](#)
- [DeleteAccessGrant](#)
- [DeleteAccessGrantsInstance](#)
- [DeleteAccessGrantsInstanceResourcePolicy](#)
- [DeleteAccessGrantsLocation](#)
- [DeleteAccessPoint](#)
- [DeleteAccessPointForObjectLambda](#)
- [DeleteAccessPointPolicy](#)
- [DeleteAccessPointPolicyForObjectLambda](#)
- [DeleteAccessPointScope](#)
- [DeleteBucket](#)
- [DeleteBucketLifecycleConfiguration](#)
- [DeleteBucketPolicy](#)
- [DeleteBucketReplication](#)
- [DeleteBucketTagging](#)
- [DeleteJobTagging](#)
- [DeleteMultiRegionAccessPoint](#)
- [DeletePublicAccessBlock](#)
- [DeleteStorageLensConfiguration](#)
- [DeleteStorageLensConfigurationTagging](#)
- [DeleteStorageLensGroup](#)

- [DescribeJob](#)
- [DescribeMultiRegionAccessPointOperation](#)
- [DissociateAccessGrantsIdentityCenter](#)
- [GetAccessGrant](#)
- [GetAccessGrantsInstance](#)
- [GetAccessGrantsInstanceForPrefix](#)
- [GetAccessGrantsInstanceResourcePolicy](#)
- [GetAccessGrantsLocation](#)
- [GetAccessPoint](#)
- [GetAccessPointConfigurationForObjectLambda](#)
- [GetAccessPointForObjectLambda](#)
- [GetAccessPointPolicy](#)
- [GetAccessPointPolicyForObjectLambda](#)
- [GetAccessPointPolicyStatus](#)
- [GetAccessPointPolicyStatusForObjectLambda](#)
- [GetAccessPointScope](#)
- [GetBucket](#)
- [GetBucketLifecycleConfiguration](#)
- [GetBucketPolicy](#)
- [GetBucketReplication](#)
- [GetBucketTagging](#)
- [GetBucketVersioning](#)
- [GetDataAccess](#)
- [GetJobTagging](#)
- [GetMultiRegionAccessPoint](#)
- [GetMultiRegionAccessPointPolicy](#)
- [GetMultiRegionAccessPointPolicyStatus](#)
- [GetMultiRegionAccessPointRoutes](#)
- [GetPublicAccessBlock](#)
- [GetStorageLensConfiguration](#)

- [GetStorageLensConfigurationTagging](#)
- [GetStorageLensGroup](#)
- [ListAccessGrants](#)
- [ListAccessGrantsInstances](#)
- [ListAccessGrantsLocations](#)
- [ListAccessPoints](#)
- [ListAccessPointsForDirectoryBuckets](#)
- [ListAccessPointsForObjectLambda](#)
- [ListCallerAccessGrants](#)
- [ListJobs](#)
- [ListMultiRegionAccessPoints](#)
- [ListRegionalBuckets](#)
- [ListStorageLensConfigurations](#)
- [ListStorageLensGroups](#)
- [ListTagsForResource](#)
- [PutAccessGrantsInstanceResourcePolicy](#)
- [PutAccessPointConfigurationForObjectLambda](#)
- [PutAccessPointPolicy](#)
- [PutAccessPointPolicyForObjectLambda](#)
- [PutAccessPointScope](#)
- [PutBucketLifecycleConfiguration](#)
- [PutBucketPolicy](#)
- [PutBucketReplication](#)
- [PutBucketTagging](#)
- [PutBucketVersioning](#)
- [PutJobTagging](#)
- [PutMultiRegionAccessPointPolicy](#)
- [PutPublicAccessBlock](#)
- [PutStorageLensConfiguration](#)
- [PutStorageLensConfigurationTagging](#)

- [SubmitMultiRegionAccessPointRoutes](#)
- [TagResource](#)
- [UntagResource](#)
- [UpdateAccessGrantsLocation](#)
- [UpdateJobPriority](#)
- [UpdateJobStatus](#)
- [UpdateStorageLensGroup](#)

The following actions are supported by Amazon S3 Files:

- [CreateAccessPoint](#)
- [CreateFileSystem](#)
- [CreateMountTarget](#)
- [DeleteAccessPoint](#)
- [DeleteFileSystem](#)
- [DeleteFileSystemPolicy](#)
- [DeleteMountTarget](#)
- [GetAccessPoint](#)
- [GetFileSystem](#)
- [GetFileSystemPolicy](#)
- [GetMountTarget](#)
- [GetSynchronizationConfiguration](#)
- [ListAccessPoints](#)
- [ListFileSystems](#)
- [ListMountTargets](#)
- [ListTagsForResource](#)
- [PutFileSystemPolicy](#)
- [PutSynchronizationConfiguration](#)
- [TagResource](#)
- [UntagResource](#)
- [UpdateMountTarget](#)

The following actions are supported by Amazon S3 on Outposts:

- [CreateEndpoint](#)
- [DeleteEndpoint](#)
- [ListEndpoints](#)
- [ListOutpostsWithS3](#)
- [ListSharedEndpoints](#)

The following actions are supported by Amazon S3 Tables:

- [CreateNamespace](#)
- [CreateTable](#)
- [CreateTableBucket](#)
- [DeleteNamespace](#)
- [DeleteTable](#)
- [DeleteTableBucket](#)
- [DeleteTableBucketEncryption](#)
- [DeleteTableBucketMetricsConfiguration](#)
- [DeleteTableBucketPolicy](#)
- [DeleteTableBucketReplication](#)
- [DeleteTablePolicy](#)
- [DeleteTableReplication](#)
- [GetNamespace](#)
- [GetTable](#)
- [GetTableBucket](#)
- [GetTableBucketEncryption](#)
- [GetTableBucketMaintenanceConfiguration](#)
- [GetTableBucketMetricsConfiguration](#)
- [GetTableBucketPolicy](#)
- [GetTableBucketReplication](#)
- [GetTableBucketStorageClass](#)
- [GetTableEncryption](#)

- [GetTableMaintenanceConfiguration](#)
- [GetTableMaintenanceJobStatus](#)
- [GetTableMetadataLocation](#)
- [GetTablePolicy](#)
- [GetTableRecordExpirationConfiguration](#)
- [GetTableRecordExpirationJobStatus](#)
- [GetTableReplication](#)
- [GetTableReplicationStatus](#)
- [GetTableStorageClass](#)
- [ListNamespaces](#)
- [ListTableBuckets](#)
- [ListTables](#)
- [ListTagsForResource](#)
- [PutTableBucketEncryption](#)
- [PutTableBucketMaintenanceConfiguration](#)
- [PutTableBucketMetricsConfiguration](#)
- [PutTableBucketPolicy](#)
- [PutTableBucketReplication](#)
- [PutTableBucketStorageClass](#)
- [PutTableMaintenanceConfiguration](#)
- [PutTablePolicy](#)
- [PutTableRecordExpirationConfiguration](#)
- [PutTableReplication](#)
- [RenameTable](#)
- [TagResource](#)
- [UntagResource](#)
- [UpdateTableMetadataLocation](#)

The following actions are supported by Amazon S3 Vectors:

- [CreateIndex](#)

- [CreateVectorBucket](#)
- [DeleteIndex](#)
- [DeleteVectorBucket](#)
- [DeleteVectorBucketPolicy](#)
- [DeleteVectors](#)
- [GetIndex](#)
- [GetVectorBucket](#)
- [GetVectorBucketPolicy](#)
- [GetVectors](#)
- [ListIndexes](#)
- [ListTagsForResource](#)
- [ListVectorBuckets](#)
- [ListVectors](#)
- [PutVectorBucketPolicy](#)
- [PutVectors](#)
- [QueryVectors](#)
- [TagResource](#)
- [UntagResource](#)

Amazon S3

The following actions are supported by Amazon S3:

- [AbortMultipartUpload](#)
- [CompleteMultipartUpload](#)
- [CopyObject](#)
- [CreateBucket](#)
- [CreateBucketMetadataConfiguration](#)
- [CreateBucketMetadataTableConfiguration](#)
- [CreateMultipartUpload](#)
- [CreateSession](#)

- [DeleteBucket](#)
- [DeleteBucketAnalyticsConfiguration](#)
- [DeleteBucketCors](#)
- [DeleteBucketEncryption](#)
- [DeleteBucketIntelligentTieringConfiguration](#)
- [DeleteBucketInventoryConfiguration](#)
- [DeleteBucketLifecycle](#)
- [DeleteBucketMetadataConfiguration](#)
- [DeleteBucketMetadataTableConfiguration](#)
- [DeleteBucketMetricsConfiguration](#)
- [DeleteBucketOwnershipControls](#)
- [DeleteBucketPolicy](#)
- [DeleteBucketReplication](#)
- [DeleteBucketTagging](#)
- [DeleteBucketWebsite](#)
- [DeleteObject](#)
- [DeleteObjectAnnotation](#)
- [DeleteObjects](#)
- [DeleteObjectTagging](#)
- [DeletePublicAccessBlock](#)
- [GetBucketAbac](#)
- [GetBucketAccelerateConfiguration](#)
- [GetBucketAcl](#)
- [GetBucketAnalyticsConfiguration](#)
- [GetBucketCors](#)
- [GetBucketEncryption](#)
- [GetBucketIntelligentTieringConfiguration](#)
- [GetBucketInventoryConfiguration](#)
- [GetBucketLifecycle](#)
- [GetBucketLifecycleConfiguration](#)

- [GetBucketLocation](#)
- [GetBucketLogging](#)
- [GetBucketMetadataConfiguration](#)
- [GetBucketMetadataTableConfiguration](#)
- [GetBucketMetricsConfiguration](#)
- [GetBucketNotification](#)
- [GetBucketNotificationConfiguration](#)
- [GetBucketOwnershipControls](#)
- [GetBucketPolicy](#)
- [GetBucketPolicyStatus](#)
- [GetBucketReplication](#)
- [GetBucketRequestPayment](#)
- [GetBucketTagging](#)
- [GetBucketVersioning](#)
- [GetBucketWebsite](#)
- [GetObject](#)
- [GetObjectAcl](#)
- [GetObjectAnnotation](#)
- [GetObjectAttributes](#)
- [GetObjectLegalHold](#)
- [GetObjectLockConfiguration](#)
- [GetObjectRetention](#)
- [GetObjectTagging](#)
- [GetObjectTorrent](#)
- [GetPublicAccessBlock](#)
- [HeadBucket](#)
- [HeadObject](#)
- [ListBucketAnalyticsConfigurations](#)
- [ListBucketIntelligentTieringConfigurations](#)
- [ListBucketInventoryConfigurations](#)

- [ListBucketMetricsConfigurations](#)
- [ListBuckets](#)
- [ListDirectoryBuckets](#)
- [ListMultipartUploads](#)
- [ListObjectAnnotations](#)
- [ListObjects](#)
- [ListObjectsV2](#)
- [ListObjectVersions](#)
- [ListParts](#)
- [PutBucketAbac](#)
- [PutBucketAccelerateConfiguration](#)
- [PutBucketAcl](#)
- [PutBucketAnalyticsConfiguration](#)
- [PutBucketCors](#)
- [PutBucketEncryption](#)
- [PutBucketIntelligentTieringConfiguration](#)
- [PutBucketInventoryConfiguration](#)
- [PutBucketLifecycle](#)
- [PutBucketLifecycleConfiguration](#)
- [PutBucketLogging](#)
- [PutBucketMetricsConfiguration](#)
- [PutBucketNotification](#)
- [PutBucketNotificationConfiguration](#)
- [PutBucketOwnershipControls](#)
- [PutBucketPolicy](#)
- [PutBucketReplication](#)
- [PutBucketRequestPayment](#)
- [PutBucketTagging](#)
- [PutBucketVersioning](#)
- [PutBucketWebsite](#)

- [PutObject](#)
- [PutObjectAcl](#)
- [PutObjectAnnotation](#)
- [PutObjectLegalHold](#)
- [PutObjectLockConfiguration](#)
- [PutObjectRetention](#)
- [PutObjectTagging](#)
- [PutPublicAccessBlock](#)
- [RenameObject](#)
- [RestoreObject](#)
- [SelectObjectContent](#)
- [UpdateBucketMetadataAnnotationTableConfiguration](#)
- [UpdateBucketMetadataInventoryTableConfiguration](#)
- [UpdateBucketMetadataJournalTableConfiguration](#)
- [UpdateObjectEncryption](#)
- [UploadPart](#)
- [UploadPartCopy](#)
- [WriteGetObjectResponse](#)

AbortMultipartUpload

Service: Amazon S3

This operation aborts a multipart upload. After a multipart upload is aborted, no additional parts can be uploaded using that upload ID. The storage consumed by any previously uploaded parts will be freed. However, if any part uploads are currently in progress, those part uploads might or might not succeed. As a result, it might be necessary to abort a given multipart upload multiple times in order to completely free all storage consumed by all parts.

To verify that all parts have been removed and prevent getting charged for the part storage, you should call the [ListParts](#) API operation and ensure that the parts list is empty.

Note

- **Directory buckets** - If multipart uploads in a directory bucket are in progress, you can't delete the bucket until all the in-progress multipart uploads are aborted or completed. To delete these in-progress multipart uploads, use the `ListMultipartUploads` operation to list the in-progress multipart uploads in the bucket and use the `AbortMultipartUpload` operation to abort all the in-progress multipart uploads.
- **Directory buckets** - For directory buckets, you must make requests for this API operation to the Zonal endpoint. These endpoints support virtual-hosted-style requests in the format `https://amzn-s3-demo-bucket.s3express-zone-id.region-code.amazonaws.com/key-name` . Path-style requests are not supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

Permissions

- **General purpose bucket permissions** - For information about permissions required to use the multipart upload, see [Multipart Upload and Permissions](#) in the *Amazon S3 User Guide*.
- **Directory bucket permissions** - To grant access to this API operation on a directory bucket, we recommend that you use the [CreateSession](#) API operation for session-based authorization. Specifically, you grant the `s3express:CreateSession` permission to the directory bucket in a bucket policy or an IAM identity-based policy. Then, you make the

CreateSession API call on the bucket to obtain a session token. With the session token in your request header, you can make API requests to this operation. After the session token expires, you make another CreateSession API call to generate a new session token for use. AWS CLI or SDKs create session and refresh the session token automatically to avoid service interruptions when a session expires. For more information about authorization, see [CreateSession](#).

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `Bucket-name.s3express-zone-id.region-code.amazonaws.com`.

The following operations are related to AbortMultipartUpload:

- [CreateMultipartUpload](#)
- [UploadPart](#)
- [CompleteMultipartUpload](#)
- [ListParts](#)
- [ListMultipartUploads](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
DELETE /Key+?uploadId=UploadId HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-request-payer: RequestPayer
x-amz-expected-bucket-owner: ExpectedBucketOwner
x-amz-if-match-initiated-time: IfMatchInitiatedTime
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name to which the upload was taking place.

Directory buckets - When you use this operation with a directory bucket, you must use virtual-hosted-style requests in the format `Bucket-name.s3express-zone-id.region-code.amazonaws.com`. Path-style requests are not supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must follow the format `bucket-base-name--zone-id--x-s3` (for example, `amzn-s3-demo-bucket--usw2-az1--x-s3`). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form `AccessPointName-AccountId.s3-accesspoint.Region.amazonaws.com`. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

Note

Object Lambda access points are not supported by directory buckets.

S3 on Outposts - When you use this action with S3 on Outposts, you must direct requests to the S3 on Outposts hostname. The S3 on Outposts hostname takes the form `AccessPointName-AccountId.outpostID.s3-outposts.Region.amazonaws.com`. When you use this action with S3 on Outposts, the destination bucket must be the Outposts access point ARN or the access point alias. For more information about S3 on Outposts, see [What is S3 on Outposts?](#) in the *Amazon S3 User Guide*.

Required: Yes

Key

Key of the object for which the multipart upload was initiated.

Length Constraints: Minimum length of 1.

Required: Yes

uploadId

Upload ID that identifies the multipart upload.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-if-match-initiated-time

If present, this header aborts an in progress multipart upload only if it was initiated on the provided timestamp. If the initiated timestamp of the multipart upload does not match the provided value, the operation returns a 412 Precondition Failed error. If the initiated timestamp matches or if the multipart upload doesn't exist, the operation returns a 204 Success (No Content) response.

Note

This functionality is only supported for directory buckets.

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
x-amz-request-charged: RequestCharged
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response.

The response returns the following HTTP headers.

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

Errors

NoSuchUpload

The specified multipart upload does not exist.

HTTP Status Code: 404

Examples

Sample Request for general purpose buckets

The following request aborts a multipart upload identified by its upload ID.

```
DELETE /example-object?
uploadId=VXBsb2FkIElEIGZvciBlbHZpbmcncyBteS1tb3ZpZS5tMnRzIHVwbG9hZ HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Mon, 1 Nov 2010 20:34:56 GMT
Authorization: authorization string
```

Sample Response for general purpose buckets

This example illustrates one usage of `AbortMultipartUpload`.

```
HTTP/1.1 204 OK
x-amz-id-2: Weag1LuByRx9e6j50nimru9p04ZVKnJ2Qz7/C1NPcfTWAtRPfTa0Fg==
x-amz-request-id: 996c76696e6727732072657175657374
Date: Mon, 1 Nov 2010 20:34:56 GMT
Content-Length: 0
Connection: keep-alive
Server: AmazonS3
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

CompleteMultipartUpload

Service: Amazon S3

Completes a multipart upload by assembling previously uploaded parts.

You first initiate the multipart upload and then upload all parts using the [UploadPart](#) operation or the [UploadPartCopy](#) operation. After successfully uploading all relevant parts of an upload, you call this CompleteMultipartUpload operation to complete the upload. Upon receiving this request, Amazon S3 concatenates all the parts in ascending order by part number to create a new object. In the CompleteMultipartUpload request, you must provide the parts list and ensure that the parts list is complete. The CompleteMultipartUpload API operation concatenates the parts that you provide in the list. For each part in the list, you must provide the PartNumber value and the ETag value that are returned after that part was uploaded.

The processing of a CompleteMultipartUpload request could take several minutes to finalize. After Amazon S3 begins processing the request, it sends an HTTP response header that specifies a 200 OK response. While processing is in progress, Amazon S3 periodically sends white space characters to keep the connection from timing out. A request could fail after the initial 200 OK response has been sent. This means that a 200 OK response can contain either a success or an error. The error response might be embedded in the 200 OK response. If you call this API operation directly, make sure to design your application to parse the contents of the response and handle it appropriately. If you use AWS SDKs, SDKs handle this condition. The SDKs detect the embedded error and apply error handling per your configuration settings (including automatically retrying the request as appropriate). If the condition persists, the SDKs throw an exception (or, for the SDKs that don't use exceptions, they return an error).

Note that if CompleteMultipartUpload fails, applications should be prepared to retry any failed requests (including 500 error responses). For more information, see [Amazon S3 Error Best Practices](#).

Important

You can't use Content-Type: application/x-www-form-urlencoded for the CompleteMultipartUpload requests. Also, if you don't provide a Content-Type header, CompleteMultipartUpload can still return a 200 OK response.

For more information about multipart uploads, see [Uploading Objects Using Multipart Upload](#) in the *Amazon S3 User Guide*.

 Note

Directory buckets - For directory buckets, you must make requests for this API operation to the Zonal endpoint. These endpoints support virtual-hosted-style requests in the format `https://amzn-s3-demo-bucket.s3express-zone-id.region-code.amazonaws.com/key-name` . Path-style requests are not supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

Permissions

- **General purpose bucket permissions** - For information about permissions required to use the multipart upload API, see [Multipart Upload and Permissions](#) in the *Amazon S3 User Guide*.

If you provide an [additional checksum value](#) in your `MultipartUpload` requests and the object is encrypted with AWS Key Management Service, you must have permission to use the `kms:Decrypt` action for the `CompleteMultipartUpload` request to succeed.

- **Directory bucket permissions** - To grant access to this API operation on a directory bucket, we recommend that you use the [CreateSession](#) API operation for session-based authorization. Specifically, you grant the `s3express:CreateSession` permission to the directory bucket in a bucket policy or an IAM identity-based policy. Then, you make the `CreateSession` API call on the bucket to obtain a session token. With the session token in your request header, you can make API requests to this operation. After the session token expires, you make another `CreateSession` API call to generate a new session token for use. AWS CLI or SDKs create session and refresh the session token automatically to avoid service interruptions when a session expires. For more information about authorization, see [CreateSession](#).

If the object is encrypted with SSE-KMS, you must also have the `kms:GenerateDataKey` and `kms:Decrypt` permissions in IAM identity-based policies and AWS KMS key policies for the AWS KMS key.

Special errors

- **Error Code: EntityTooSmall**
 - **Description:** Your proposed upload is smaller than the minimum allowed object size. Each part must be at least 5 MB in size, except the last part.
 - **HTTP Status Code:** 400 Bad Request
- **Error Code: InvalidPart**
 - **Description:** One or more of the specified parts could not be found. The part might not have been uploaded, or the specified ETag might not have matched the uploaded part's ETag.
 - **HTTP Status Code:** 400 Bad Request
- **Error Code: InvalidPartOrder**
 - **Description:** The list of parts was not in ascending order. The parts list must be specified in order by part number.
 - **HTTP Status Code:** 400 Bad Request
- **Error Code: NoSuchUpload**
 - **Description:** The specified multipart upload does not exist. The upload ID might be invalid, or the multipart upload might have been aborted or completed.
 - **HTTP Status Code:** 404 Not Found

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `Bucket-name.s3express-zone-id.region-code.amazonaws.com`.

The following operations are related to `CompleteMultipartUpload`:

- [CreateMultipartUpload](#)
- [UploadPart](#)
- [AbortMultipartUpload](#)
- [ListParts](#)
- [ListMultipartUploads](#)

⚠ Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
POST /Key+?uploadId=UploadId HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-checksum-crc32: ChecksumCRC32
x-amz-checksum-crc32c: ChecksumCRC32C
x-amz-checksum-crc64nvme: ChecksumCRC64NVME
x-amz-checksum-sha1: ChecksumSHA1
x-amz-checksum-sha256: ChecksumSHA256
x-amz-checksum-sha512: ChecksumSHA512
x-amz-checksum-md5: ChecksumMD5
x-amz-checksum-xxhash64: ChecksumXXHASH64
x-amz-checksum-xxhash3: ChecksumXXHASH3
x-amz-checksum-xxhash128: ChecksumXXHASH128
x-amz-checksum-type: ChecksumType
x-amz-mp-object-size: MpuObjectSize
x-amz-request-payer: RequestPayer
x-amz-expected-bucket-owner: ExpectedBucketOwner
If-Match: IfMatch
If-None-Match: IfNoneMatch
x-amz-server-side-encryption-customer-algorithm: SSECustomerAlgorithm
x-amz-server-side-encryption-customer-key: SSECustomerKey
x-amz-server-side-encryption-customer-key-MD5: SSECustomerKeyMD5
<?xml version="1.0" encoding="UTF-8"?>
<CompleteMultipartUpload xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Part>
    <ChecksumCRC32>string</ChecksumCRC32>
    <ChecksumCRC32C>string</ChecksumCRC32C>
    <ChecksumCRC64NVME>string</ChecksumCRC64NVME>
    <ChecksumMD5>string</ChecksumMD5>
    <ChecksumSHA1>string</ChecksumSHA1>
    <ChecksumSHA256>string</ChecksumSHA256>
    <ChecksumSHA512>string</ChecksumSHA512>
    <ChecksumXXHASH128>string</ChecksumXXHASH128>
    <ChecksumXXHASH3>string</ChecksumXXHASH3>
```



```
<ChecksumXXHASH64>string</ChecksumXXHASH64>
<ETag>string</ETag>
<PartNumber>integer</PartNumber>
</Part>
...
</CompleteMultipartUpload>
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

Name of the bucket to which the multipart upload was initiated.

Directory buckets - When you use this operation with a directory bucket, you must use virtual-hosted-style requests in the format *Bucket-name.s3express-zone-id.region-code.amazonaws.com*. Path-style requests are not supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must follow the format *bucket-base-name--zone-id--x-s3* (for example, *amzn-s3-demo-bucket--usw2-az1--x-s3*). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form *AccessPointName-AccountId.s3-accesspoint.Region.amazonaws.com*. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

Note

Object Lambda access points are not supported by directory buckets.

S3 on Outposts - When you use this action with S3 on Outposts, you must direct requests to the S3 on Outposts hostname. The S3 on Outposts hostname takes the form

`AccessPointName-AccountId.outpostID.s3-outposts.Region.amazonaws.com`.
When you use this action with S3 on Outposts, the destination bucket must be the Outposts access point ARN or the access point alias. For more information about S3 on Outposts, see [What is S3 on Outposts?](#) in the *Amazon S3 User Guide*.

Required: Yes

If-Match

Uploads the object only if the ETag (entity tag) value provided during the WRITE operation matches the ETag of the object in S3. If the ETag values do not match, the operation returns a 412 Precondition Failed error.

If a conflicting operation occurs during the upload S3 returns a 409 ConditionalRequestConflict response. On a 409 failure you should fetch the object's ETag, re-initiate the multipart upload with `CreateMultipartUpload`, and re-upload each part.

Expects the ETag value as a string.

For more information about conditional requests, see [RFC 7232](#), or [Conditional requests](#) in the *Amazon S3 User Guide*.

If-None-Match

Uploads the object only if the object key name does not already exist in the bucket specified. Otherwise, Amazon S3 returns a 412 Precondition Failed error.

If a conflicting operation occurs during the upload S3 returns a 409 ConditionalRequestConflict response. On a 409 failure you should re-initiate the multipart upload with `CreateMultipartUpload` and re-upload each part.

Expects the '*' (asterisk) character.

For more information about conditional requests, see [RFC 7232](#), or [Conditional requests](#) in the *Amazon S3 User Guide*.

Key

Object key for which the multipart upload was initiated.

Length Constraints: Minimum length of 1.

Required: Yes

uploadId

ID for the initiated multipart upload.

Required: Yes

x-amz-checksum-crc32

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 32-bit CRC32 checksum of the object. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-crc32c

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 32-bit CRC32C checksum of the object. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-crc64nvme

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 64-bit CRC64NVME checksum of the object. The CRC64NVME checksum is always a full object checksum. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-checksum-md5

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 128-bit MD5 digest of the object. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-checksum-sha1

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 160-bit SHA1 digest of the object. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-sha256

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 256-bit SHA256 digest of the object. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-sha512

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 512-bit SHA512 digest of the object. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-checksum-type

This header specifies the checksum type of the object, which determines how part-level checksums are combined to create an object-level checksum for multipart objects. You can use this header as a data integrity check to verify that the checksum type that is received is the same checksum that was specified. If the checksum type doesn't match the checksum type that was specified for the object during the CreateMultipartUpload request, it'll result in a BadDigest error. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

Valid Values: COMPOSITE | FULL_OBJECT

x-amz-checksum-xxhash128

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 128-bit XXHASH128 checksum of the object. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-checksum-xxhash3

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 64-bit XXHASH3 checksum of the object. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-checksum-xxhash64

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 64-bit XXHASH64 checksum of the object. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-mp-object-size

The expected total object size of the multipart upload request. If there's a mismatch between the specified object size value and the actual object size value, it results in an HTTP 400 InvalidRequest error.

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

x-amz-server-side-encryption-customer-algorithm

The server-side encryption (SSE) algorithm used to encrypt the object. This parameter is required only when the object was created using a checksum algorithm or if your bucket policy requires the use of SSE-C. For more information, see [Protecting data using SSE-C keys](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

x-amz-server-side-encryption-customer-key

The server-side encryption (SSE) customer managed key. This parameter is needed only when the object was created using a checksum algorithm. For more information, see [Protecting data using SSE-C keys](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

x-amz-server-side-encryption-customer-key-MD5

The MD5 server-side encryption (SSE) customer managed key. This parameter is needed only when the object was created using a checksum algorithm. For more information, see [Protecting data using SSE-C keys](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Request Body

The request accepts the following data in XML format.

CompleteMultipartUpload

Root level tag for the CompleteMultipartUpload parameters.

Required: Yes

Part

Array of CompletedPart data types.

If you do not supply a valid Part with your request, the service sends back an HTTP 400 response.

Type: Array of [CompletedPart](#) data types

Required: No

Response Syntax

HTTP/1.1 200

```

x-amz-expiration: Expiration
x-amz-server-side-encryption: ServerSideEncryption
x-amz-version-id: VersionId
x-amz-server-side-encryption-aws-kms-key-id: SSEKMSKeyId
x-amz-server-side-encryption-bucket-key-enabled: BucketKeyEnabled
x-amz-request-charged: RequestCharged
<?xml version="1.0" encoding="UTF-8"?>
<CompleteMultipartUploadResult>
  <Location>string</Location>
  <Bucket>string</Bucket>
  <Key>string</Key>
  <ETag>string</ETag>
  <ChecksumCRC32>string</ChecksumCRC32>
  <ChecksumCRC32C>string</ChecksumCRC32C>
  <ChecksumCRC64NVME>string</ChecksumCRC64NVME>
  <ChecksumSHA1>string</ChecksumSHA1>
  <ChecksumSHA256>string</ChecksumSHA256>
  <ChecksumSHA512>string</ChecksumSHA512>
  <ChecksumMD5>string</ChecksumMD5>
  <ChecksumXXHASH64>string</ChecksumXXHASH64>
  <ChecksumXXHASH3>string</ChecksumXXHASH3>
  <ChecksumXXHASH128>string</ChecksumXXHASH128>
  <ChecksumType>string</ChecksumType>
</CompleteMultipartUploadResult>

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

x-amz-expiration

If the object expiration is configured, this will contain the expiration date (expiry-date) and rule ID (rule-id). The value of rule-id is URL-encoded.

Note

This functionality is not supported for directory buckets.

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

x-amz-server-side-encryption

The server-side encryption algorithm used when storing this object in Amazon S3.

Note

When accessing data stored in Amazon FSx file systems using S3 access points, the only valid server side encryption option is `aws:fsx`.

Valid Values: AES256 | `aws:fsx` | `aws:backup` | `aws:kms` | `aws:kms:dsse`

x-amz-server-side-encryption-aws-kms-key-id

If present, indicates the ID of the KMS key that was used for object encryption.

x-amz-server-side-encryption-bucket-key-enabled

Indicates whether the multipart upload uses an S3 Bucket Key for server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS).

x-amz-version-id

Version ID of the newly created object, in case the bucket has versioning turned on.

Note

This functionality is not supported for directory buckets.

The following data is returned in XML format by the service.

CompleteMultipartUploadResult

Root level tag for the CompleteMultipartUploadResult parameters.

Required: Yes

Bucket

The name of the bucket that contains the newly created object. Does not return the access point ARN or access point alias if used.

Note

Access points are not supported by directory buckets.

Type: String

ChecksumCRC32

The Base64 encoded, 32-bit CRC32 checksum of the object. This checksum is only present if the checksum was uploaded with the object. When you use an API operation on an object that was uploaded using multipart uploads, this value may not be a direct checksum value of the full object. Instead, it's a calculation based on the checksum values of each individual part. For more information about how checksums are calculated with multipart uploads, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

ChecksumCRC32C

The Base64 encoded, 32-bit CRC32C checksum of the object. This checksum is only present if the checksum was uploaded with the object. When you use an API operation on an object that was uploaded using multipart uploads, this value may not be a direct checksum value of the full object. Instead, it's a calculation based on the checksum values of each individual part. For more information about how checksums are calculated with multipart uploads, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

ChecksumCRC64NVME

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 64-bit CRC64NVME checksum of the object. The CRC64NVME checksum is always a full object checksum. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

Type: String

ChecksumMD5

The Base64 encoded, 128-bit MD5 digest of the object. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

Type: String

ChecksumSHA1

The Base64 encoded, 160-bit SHA1 digest of the object. This checksum is only present if the checksum was uploaded with the object. When you use the API operation on an object that was uploaded using multipart uploads, this value may not be a direct checksum value of the full object. Instead, it's a calculation based on the checksum values of each individual part. For more information about how checksums are calculated with multipart uploads, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

ChecksumSHA256

The Base64 encoded, 256-bit SHA256 digest of the object. This checksum is only present if the checksum was uploaded with the object. When you use an API operation on an object that was uploaded using multipart uploads, this value may not be a direct checksum value of the full object. Instead, it's a calculation based on the checksum values of each individual part. For more information about how checksums are calculated with multipart uploads, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

ChecksumSHA512

The Base64 encoded, 512-bit SHA512 digest of the object. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

Type: String

ChecksumType

The checksum type, which determines how part-level checksums are combined to create an object-level checksum for multipart objects. You can use this header as a data integrity check to verify that the checksum type that is received is the same checksum type that was specified during the `CreateMultipartUpload` request. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

Type: String

Valid Values: COMPOSITE | FULL_OBJECT

ChecksumXXHASH128

The Base64 encoded, 128-bit XXHASH128 checksum of the object. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

Type: String

ChecksumXXHASH3

The Base64 encoded, 64-bit XXHASH3 checksum of the object. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

Type: String

ChecksumXXHASH64

The Base64 encoded, 64-bit XXHASH64 checksum of the object. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

Type: String

ETag

Entity tag that identifies the newly created object's data. Objects with different object data will have different entity tags. The entity tag is an opaque string. The entity tag may or may not be an MD5 digest of the object data. If the entity tag is not an MD5 digest of the object data, it will contain one or more nonhexadecimal characters and/or will consist of less than 32 or more than 32 hexadecimal digits. For more information about how the entity tag is calculated, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Key

The object key of the newly created object.

Type: String

Length Constraints: Minimum length of 1.

Location

The URI that identifies the newly created object.

Type: String

Examples

Sample Request for general purpose buckets

The following Complete Multipart Upload request specifies three parts in the CompleteMultipartUpload element.

```
POST /example-object?
uploadId=AAAsb2FkIElEIGZvciBlbHZpbmcncyWeeS1tb3ZpZS5tMnRzIRRwbG9hZA HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Mon, 1 Nov 2010 20:34:56 GMT
Content-Length: 391
Authorization: authorization string

<CompleteMultipartUpload>
  <Part>
    <PartNumber>1</PartNumber>
    <ETag>"a54357aff0632cce46d942af68356b38"</ETag>
  </Part>
  <Part>
    <PartNumber>2</PartNumber>
    <ETag>"0c78aef83f66abc1fa1e8477f296d394"</ETag>
  </Part>
  <Part>
    <PartNumber>3</PartNumber>
    <ETag>"acbd18db4cc2f85cedef654fccc4a4d8"</ETag>
  </Part>
</CompleteMultipartUpload>
```

Sample Response for general purpose buckets

The following response indicates that an object was successfully assembled.

```
HTTP/1.1 200 OK
x-amz-id-2: Uuag1LuByRx9e6j50nimru9p04ZVKnJ2Qz7/C1NPcfTWAtRPfTa0Fg==
x-amz-request-id: 656c76696e6727732072657175657374
Date: Mon, 1 Nov 2010 20:34:56 GMT
Connection: close
Server: AmazonS3

<?xml version="1.0" encoding="UTF-8"?>
<CompleteMultipartUploadResult xmlns="http://s3.amazonaws.com/
doc/2006-03-01/">
  <Location>http://amzn-s3-demo-bucket.s3.<Region>.amazonaws.com/Example-
Object</Location>
  <Bucket>amzn-s3-demo-bucket</Bucket>
  <Key>Example-Object</Key>
  <ETag>"3858f62230ac3c915f300c664312c11f-9"</ETag>
</CompleteMultipartUploadResult>
```

Sample Response for general purpose buckets: Error specified in header

The following response indicates that an error occurred before the HTTP response header was sent.

```
HTTP/1.1 403 Forbidden
x-amz-id-2: Uuag1LuByRx9e6j50nimru9p04ZVKnJ2Qz7/C1NPcfTWAtRPfTa0Fg==
x-amz-request-id: 656c76696e6727732072657175657374
Date: Mon, 1 Nov 2010 20:34:56 GMT
Content-Length: 237
Connection: keep-alive
Server: AmazonS3

<?xml version="1.0" encoding="UTF-8"?>
<Error>
  <Code>AccessDenied</Code>
  <Message>Access Denied</Message>
  <RequestId>656c76696e6727732072657175657374</RequestId>
```

```
<HostId>Uuag1LuByRx9e6j50nimru9p04ZVKnJ2Qz7/C1NPcfTWAtRPfTa0Fg==</HostId>  
</Error>
```

Sample Response for general purpose buckets: Error specified in body

The following response indicates that an error occurred after the HTTP response header was sent. Note that while the HTTP status code is 200 OK, the request actually failed as described in the `Error` element.

```
HTTP/1.1 200 OK  
x-amz-id-2: Uuag1LuByRx9e6j50nimru9p04ZVKnJ2Qz7/C1NPcfTWAtRPfTa0Fg==  
x-amz-request-id: 656c76696e6727732072657175657374  
Date: Mon, 1 Nov 2010 20:34:56 GMT  
Connection: close  
Server: AmazonS3  
  
<?xml version="1.0" encoding="UTF-8"?>  
  
<Error>  
  <Code>InternalError</Code>  
  <Message>We encountered an internal error. Please try again.</Message>  
  <RequestId>656c76696e6727732072657175657374</RequestId>  
  <HostId>Uuag1LuByRx9e6j50nimru9p04ZVKnJ2Qz7/C1NPcfTWAtRPfTa0Fg==</HostId>  
</Error>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)

- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

CopyObject

Service: Amazon S3

Creates a copy of an object that is already stored in Amazon S3.

Important

End of support notice: As of October 1, 2025, Amazon S3 has discontinued support for Email Grantee Access Control Lists (ACLs). If you attempt to use an Email Grantee ACL in a request after October 1, 2025, the request will receive an HTTP 405 (Method Not Allowed) error.

This change affects the following AWS Regions: US East (N. Virginia), US West (N. California), US West (Oregon), Asia Pacific (Singapore), Asia Pacific (Sydney), Asia Pacific (Tokyo), Europe (Ireland), and South America (São Paulo).

Note

You can store individual objects of up to 50 TB in Amazon S3. You create a copy of your object up to 5 GB in size in a single atomic action using this API. However, to copy an object greater than 5 GB, you must use the multipart upload Upload Part - Copy (UploadPartCopy) API. For more information, see [Copy Object Using the REST Multipart Upload API](#).

You can copy individual objects between general purpose buckets, between directory buckets, and between general purpose buckets and directory buckets.

Note

- Amazon S3 supports copy operations using Multi-Region Access Points only as a destination when using the Multi-Region Access Point ARN.
- **Directory buckets** - For directory buckets, you must make requests for this API operation to the Zonal endpoint. These endpoints support virtual-hosted-style requests in the format `https://amzn-s3-demo-bucket.s3express-zone-id.region-code.amazonaws.com/key-name` . Path-style requests are not supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints](#)

[for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

- VPC endpoints don't support cross-Region requests (including copies). If you're using VPC endpoints, your source and destination buckets should be in the same AWS Region as your VPC endpoint.

Both the Region that you want to copy the object from and the Region that you want to copy the object to must be enabled for your account. For more information about how to enable a Region for your account, see [Enable or disable a Region for standalone accounts](#) in the *AWS Account Management Guide*.

Important

Amazon S3 transfer acceleration does not support cross-Region copies. If you request a cross-Region copy using a transfer acceleration endpoint, you get a 400 Bad Request error. For more information, see [Transfer Acceleration](#).

Authentication and authorization

All CopyObject requests must be authenticated and signed by using IAM credentials (access key ID and secret access key for the IAM identities). All headers with the x-amz- prefix, including x-amz-copy-source, must be signed. For more information, see [REST Authentication](#).

Directory buckets - You must use the IAM credentials to authenticate and authorize your access to the CopyObject API operation, instead of using the temporary security credentials through the CreateSession API operation.

AWS CLI or SDKs handles authentication and authorization on your behalf.

Permissions

You must have *read* access to the source object and *write* access to the destination bucket.

- **General purpose bucket permissions** - You must have permissions in an IAM policy based on the source and destination bucket types in a CopyObject operation.

- If the source object is in a general purpose bucket, you must have **s3:GetObject** permission to read the source object that is being copied.
- If the destination bucket is a general purpose bucket, you must have **s3:PutObject** permission to write the object copy to the destination bucket.
- **Directory bucket permissions** - You must have permissions in a bucket policy or an IAM identity-based policy based on the source and destination bucket types in a CopyObject operation.
 - If the source object that you want to copy is in a directory bucket, you must have the **s3express:CreateSession** permission in the Action element of a policy to read the object. If no session mode is specified, the session will be created with the maximum allowable privilege, attempting ReadWrite first, then ReadOnly if ReadWrite is not permitted. If you want to explicitly restrict the access to be read-only, you can set the **s3express:SessionMode** condition key to ReadOnly on the copy source bucket.
 - If the copy destination is a directory bucket, you must have the **s3express:CreateSession** permission in the Action element of a policy to write the object to the destination. The **s3express:SessionMode** condition key can't be set to ReadOnly on the copy destination bucket.

If the object is encrypted with SSE-KMS, you must also have the **kms:GenerateDataKey** and **kms:Decrypt** permissions in IAM identity-based policies and AWS KMS key policies for the AWS KMS key.

For example policies, see [Example bucket policies for S3 Express One Zone](#) and [AWS Identity and Access Management \(IAM\) identity-based policies for S3 Express One Zone](#) in the *Amazon S3 User Guide*.

Response and special errors

When the request is an HTTP 1.1 request, the response is chunk encoded. When the request is not an HTTP 1.1 request, the response would not contain the Content-Length. You always need to read the entire response body to check if the copy succeeds.

- If the copy is successful, you receive a response with information about the copied object.
- A copy request might return an error when Amazon S3 receives the copy request or while Amazon S3 is copying the files. A 200 OK response can contain either a success or an error.
- If the error occurs before the copy action starts, you receive a standard Amazon S3 error.

- If the error occurs during the copy operation, the error response is embedded in the 200 OK response. For example, in a cross-region copy, you may encounter throttling and receive a 200 OK response. For more information, see [Resolve the Error 200 response when copying objects to Amazon S3](#). The 200 OK status code means the copy was accepted, but it doesn't mean the copy is complete. Another example is when you disconnect from Amazon S3 before the copy is complete, Amazon S3 might cancel the copy and you may receive a 200 OK response. You must stay connected to Amazon S3 until the entire response is successfully received and processed.

If you call this API operation directly, make sure to design your application to parse the content of the response and handle it appropriately. If you use AWS SDKs, SDKs handle this condition. The SDKs detect the embedded error and apply error handling per your configuration settings (including automatically retrying the request as appropriate). If the condition persists, the SDKs throw an exception (or, for the SDKs that don't use exceptions, they return an error).

Charge

The copy request charge is based on the storage class and Region that you specify for the destination object. The request can also result in a data retrieval charge for the source if the source storage class bills for data retrieval. If the copy source is in a different region, the data transfer is billed to the copy source account. For pricing information, see [Amazon S3 pricing](#).

HTTP Host header syntax

- **Directory buckets** - The HTTP Host header syntax is `Bucket-name.s3express-zone-id.region-code.amazonaws.com`.
- **Amazon S3 on Outposts** - When you use this action with S3 on Outposts through the REST API, you must direct requests to the S3 on Outposts hostname. The S3 on Outposts hostname takes the form `AccessPointName-AccountId.outpostID.s3-outposts.Region.amazonaws.com`. The hostname isn't required when you use the AWS CLI or SDKs.

The following operations are related to CopyObject:

- [PutObject](#)
- [GetObject](#)

⚠ Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /Key+ HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-acl: ACL
Cache-Control: CacheControl
x-amz-checksum-algorithm: ChecksumAlgorithm
Content-Disposition: ContentDisposition
Content-Encoding: ContentEncoding
Content-Language: ContentLanguage
Content-Type: ContentType
x-amz-copy-source: CopySource
x-amz-copy-source-if-match: CopySourceIfMatch
x-amz-copy-source-if-modified-since: CopySourceIfModifiedSince
x-amz-copy-source-if-none-match: CopySourceIfNoneMatch
x-amz-copy-source-if-unmodified-since: CopySourceIfUnmodifiedSince
Expires: Expires
x-amz-grant-full-control: GrantFullControl
x-amz-grant-read: GrantRead
x-amz-grant-read-acp: GrantReadACP
x-amz-grant-write-acp: GrantWriteACP
If-Match: IfMatch
If-None-Match: IfNoneMatch
x-amz-metadata-directive: MetadataDirective
x-amz-tagging-directive: TaggingDirective
x-amz-object-annotation-directive: AnnotationDirective
x-amz-server-side-encryption: ServerSideEncryption
x-amz-storage-class: StorageClass
x-amz-website-redirect-location: WebsiteRedirectLocation
x-amz-server-side-encryption-customer-algorithm: SSECustomerAlgorithm
x-amz-server-side-encryption-customer-key: SSECustomerKey
x-amz-server-side-encryption-customer-key-MD5: SSECustomerKeyMD5
x-amz-server-side-encryption-aws-kms-key-id: SSEKMSKeyId
x-amz-server-side-encryption-context: SSEKMSEncryptionContext
x-amz-server-side-encryption-bucket-key-enabled: BucketKeyEnabled
```

```

x-amz-copy-source-server-side-encryption-customer-
algorithm: CopySourceSSECustomerAlgorithm
x-amz-copy-source-server-side-encryption-customer-key: CopySourceSSECustomerKey
x-amz-copy-source-server-side-encryption-customer-key-MD5: CopySourceSSECustomerKeyMD5
x-amz-request-payer: RequestPayer
x-amz-tagging: Tagging
x-amz-object-lock-mode: ObjectLockMode
x-amz-object-lock-retain-until-date: ObjectLockRetainUntilDate
x-amz-object-lock-legal-hold: ObjectLockLegalHoldStatus
x-amz-object-lock-event-hold: ObjectLockEventHold
x-amz-object-lock-event-hold-duration-days: ObjectLockEventHoldDurationDays
x-amz-object-lock-event-hold-duration-years: ObjectLockEventHoldDurationYears
x-amz-expected-bucket-owner: ExpectedBucketOwner
x-amz-source-expected-bucket-owner: ExpectedSourceBucketOwner

```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the destination bucket.

Directory buckets - When you use this operation with a directory bucket, you must use virtual-hosted-style requests in the format *Bucket-name.s3express-zone-id.region-code.amazonaws.com*. Path-style requests are not supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must follow the format *bucket-base-name--zone-id--x-s3* (for example, *amzn-s3-demo-bucket--usw2-az1--x-s3*). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*.

Note

Copying objects across different AWS Regions isn't supported when the source or destination bucket is in AWS Local Zones. The source and destination buckets must have the same parent AWS Region. Otherwise, you get an HTTP 400 Bad Request error with the error code `InvalidRequest`.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access

point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form *AccessPointName-AccountId.s3-accesspoint.Region.amazonaws.com*. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

 Note

Object Lambda access points are not supported by directory buckets.

S3 on Outposts - When you use this action with S3 on Outposts, you must use the Outpost bucket access point ARN or the access point alias for the destination bucket. You can only copy objects within the same Outpost bucket. It's not supported to copy objects across different AWS Outposts, between buckets on the same Outposts, or between Outposts buckets and any other bucket types. For more information about S3 on Outposts, see [What is S3 on Outposts?](#) in the *S3 on Outposts guide*. When you use this action with S3 on Outposts through the REST API, you must direct requests to the S3 on Outposts hostname, in the format *AccessPointName-AccountId.outpostID.s3-outposts.Region.amazonaws.com*. The hostname isn't required when you use the AWS CLI or SDKs.

Required: Yes

Cache-Control

Specifies the caching behavior along the request/reply chain.

Content-Disposition

Specifies presentational information for the object. Indicates whether an object should be displayed in a web browser or downloaded as a file. It allows specifying the desired filename for the downloaded file.

Content-Encoding

Specifies what content encodings have been applied to the object and thus what decoding mechanisms must be applied to obtain the media-type referenced by the Content-Type header field.

Note

For directory buckets, only the `aws-chunked` value is supported in this header field.

Content-Language

The language the content is in.

Content-Type

A standard MIME type that describes the format of the object data.

Expires

The date and time at which the object is no longer cacheable.

If-Match

Copies the object if the entity tag (ETag) of the destination object matches the specified tag. If the ETag values do not match, the operation returns a `412 Precondition Failed` error. If a concurrent operation occurs during the upload S3 returns a `409 ConditionalRequestConflict` response. On a 409 failure you should fetch the object's ETag and retry the upload.

Expects the ETag value as a string.

For more information about conditional requests, see [RFC 7232](#).

If-None-Match

Copies the object only if the object key name at the destination does not already exist in the bucket specified. Otherwise, Amazon S3 returns a `412 Precondition Failed` error. If a concurrent operation occurs during the upload S3 returns a `409 ConditionalRequestConflict` response. On a 409 failure you should retry the upload.

Expects the `'*'` (asterisk) character.

For more information about conditional requests, see [RFC 7232](#).

Key

The key of the destination object.

Length Constraints: Minimum length of 1.

Required: Yes

[x-amz-acl](#)

The canned access control list (ACL) to apply to the object.

When you copy an object, the ACL metadata is not preserved and is set to `private` by default. Only the owner has full access control. To override the default ACL setting, specify a new ACL when you generate a copy request. For more information, see [Using ACLs](#).

If the destination bucket that you're copying objects to uses the bucket owner enforced setting for S3 Object Ownership, ACLs are disabled and no longer affect permissions. Buckets that use this setting only accept PUT requests that don't specify an ACL or PUT requests that specify bucket owner full control ACLs, such as the `bucket-owner-full-control` canned ACL or an equivalent form of this ACL expressed in the XML format. For more information, see [Controlling ownership of objects and disabling ACLs](#) in the *Amazon S3 User Guide*.

Note

- If your destination bucket uses the bucket owner enforced setting for Object Ownership, all objects written to the bucket by any account will be owned by the bucket owner.
- This functionality is not supported for directory buckets.
- This functionality is not supported for Amazon S3 on Outposts.

Valid Values: `private` | `public-read` | `public-read-write` | `authenticated-read` | `aws-exec-read` | `bucket-owner-read` | `bucket-owner-full-control`

[x-amz-checksum-algorithm](#)

Indicates the algorithm that you want Amazon S3 to use to create the checksum for the object. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

When you copy an object, if the source object has a checksum, that checksum value will be copied to the new object by default. If the `CopyObject` request does not include this `x-amz-checksum-algorithm` header, the checksum algorithm will be copied from the source object to the destination object (if it's present on the source object). You can optionally specify a different checksum algorithm to use with the `x-amz-checksum-algorithm`

header. Unrecognized or unsupported values will respond with the HTTP status code 400 Bad Request.

Note

For directory buckets, when you use AWS SDKs, CRC32 is the default checksum algorithm that's used for performance.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 | XXHASH64 | XXHASH3 | XXHASH128

x-amz-copy-source

Specifies the source object for the copy operation. The source object can be up to 5 GB. If the source object is an object that was uploaded by using a multipart upload, the object copy will be a single part object after the source object is copied to the destination bucket.

You specify the value of the copy source in one of two formats, depending on whether you want to access the source object through an [access point](#):

- For objects not accessed through an access point, specify the name of the source bucket and the key of the source object, separated by a slash (/). For example, to copy the object `reports/january.pdf` from the general purpose bucket `awsexamplebucket`, use `awsexamplebucket/reports/january.pdf`. The value must be URL-encoded. To copy the object `reports/january.pdf` from the directory bucket `awsexamplebucket--use1-az5--x-s3`, use `awsexamplebucket--use1-az5--x-s3/reports/january.pdf`. The value must be URL-encoded.
- For objects accessed through access points, specify the Amazon Resource Name (ARN) of the object as accessed through the access point, in the format `arn:aws:s3:<Region>:<account-id>:accesspoint/<access-point-name>/object/<key>`. For example, to copy the object `reports/january.pdf` through access point `my-access-point` owned by account `123456789012` in Region `us-west-2`, use the URL encoding of `arn:aws:s3:us-west-2:123456789012:accesspoint/my-access-point/object/reports/january.pdf`. The value must be URL encoded.

Note

- Amazon S3 supports copy operations using Access points only when the source and destination buckets are in the same AWS Region.

- Access points are not supported by directory buckets.

Alternatively, for objects accessed through Amazon S3 on Outposts, specify the ARN of the object as accessed in the format `arn:aws:s3-outposts:<Region>:<account-id>:outpost/<outpost-id>/object/<key>`. For example, to copy the object `reports/january.pdf` through outpost `my-outpost` owned by account `123456789012` in Region `us-west-2`, use the URL encoding of `arn:aws:s3-outposts:us-west-2:123456789012:outpost/my-outpost/object/reports/january.pdf`. The value must be URL-encoded.

If your source bucket versioning is enabled, the `x-amz-copy-source` header by default identifies the current version of an object to copy. If the current version is a delete marker, Amazon S3 behaves as if the object was deleted. To copy a different version, use the `versionId` query parameter. Specifically, append `?versionId=<version-id>` to the value (for example, `awsexamplebucket/reports/january.pdf?versionId=QUpfndndhfd8438MNFdN93jdnJFkdmqnh893`). If you don't specify a version ID, Amazon S3 copies the latest version of the source object.

If you enable versioning on the destination bucket, Amazon S3 generates a unique version ID for the copied object. This version ID is different from the version ID of the source object. Amazon S3 returns the version ID of the copied object in the `x-amz-version-id` response header in the response.

If you do not enable versioning or suspend it on the destination bucket, the version ID that Amazon S3 generates in the `x-amz-version-id` response header is always null.

 Note

Directory buckets - S3 Versioning isn't enabled and supported for directory buckets.

Pattern: `\/?\.\+\/\.\+`

Required: Yes

[x-amz-copy-source-if-match](#)

Copies the object if its entity tag (ETag) matches the specified tag.

If both the `x-amz-copy-source-if-match` and `x-amz-copy-source-if-unmodified-since` headers are present in the request and evaluate as follows, Amazon S3 returns 200 OK and copies the data:

- `x-amz-copy-source-if-match` condition evaluates to true
- `x-amz-copy-source-if-unmodified-since` condition evaluates to false

[x-amz-copy-source-if-modified-since](#)

Copies the object if it has been modified since the specified time.

If both the `x-amz-copy-source-if-none-match` and `x-amz-copy-source-if-modified-since` headers are present in the request and evaluate as follows, Amazon S3 returns the 412 Precondition Failed response code:

- `x-amz-copy-source-if-none-match` condition evaluates to false
- `x-amz-copy-source-if-modified-since` condition evaluates to true

[x-amz-copy-source-if-none-match](#)

Copies the object if its entity tag (ETag) is different than the specified ETag.

If both the `x-amz-copy-source-if-none-match` and `x-amz-copy-source-if-modified-since` headers are present in the request and evaluate as follows, Amazon S3 returns the 412 Precondition Failed response code:

- `x-amz-copy-source-if-none-match` condition evaluates to false
- `x-amz-copy-source-if-modified-since` condition evaluates to true

[x-amz-copy-source-if-unmodified-since](#)

Copies the object if it hasn't been modified since the specified time.

If both the `x-amz-copy-source-if-match` and `x-amz-copy-source-if-unmodified-since` headers are present in the request and evaluate as follows, Amazon S3 returns 200 OK and copies the data:

- `x-amz-copy-source-if-match` condition evaluates to true
- `x-amz-copy-source-if-unmodified-since` condition evaluates to false

[x-amz-copy-source-server-side-encryption-customer-algorithm](#)

Specifies the algorithm to use when decrypting the source object (for example, AES256).

If the source object for the copy is stored in Amazon S3 using SSE-C, you must provide the necessary encryption information in your request so that Amazon S3 can decrypt the object for copying.

 Note

This functionality is not supported when the source object is in a directory bucket.

x-amz-copy-source-server-side-encryption-customer-key

Specifies the customer-provided encryption key for Amazon S3 to use to decrypt the source object. The encryption key provided in this header must be the same one that was used when the source object was created.

If the source object for the copy is stored in Amazon S3 using SSE-C, you must provide the necessary encryption information in your request so that Amazon S3 can decrypt the object for copying.

 Note

This functionality is not supported when the source object is in a directory bucket.

x-amz-copy-source-server-side-encryption-customer-key-MD5

Specifies the 128-bit MD5 digest of the encryption key according to RFC 1321. Amazon S3 uses this header for a message integrity check to ensure that the encryption key was transmitted without error.

If the source object for the copy is stored in Amazon S3 using SSE-C, you must provide the necessary encryption information in your request so that Amazon S3 can decrypt the object for copying.

 Note

This functionality is not supported when the source object is in a directory bucket.

x-amz-expected-bucket-owner

The account ID of the expected destination bucket owner. If the account ID that you provide does not match the actual owner of the destination bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-grant-full-control

Gives the grantee READ, READ_ACP, and WRITE_ACP permissions on the object.

Note

- This functionality is not supported for directory buckets.
- This functionality is not supported for Amazon S3 on Outposts.

x-amz-grant-read

Allows grantee to read the object data and its metadata.

Note

- This functionality is not supported for directory buckets.
- This functionality is not supported for Amazon S3 on Outposts.

x-amz-grant-read-acp

Allows grantee to read the object ACL.

Note

- This functionality is not supported for directory buckets.
- This functionality is not supported for Amazon S3 on Outposts.

x-amz-grant-write-acp

Allows grantee to write the ACL for the applicable object.

Note

- This functionality is not supported for directory buckets.
- This functionality is not supported for Amazon S3 on Outposts.

x-amz-metadata-directive

Specifies whether the metadata is copied from the source object or replaced with metadata that's provided in the request. When copying an object, you can preserve all metadata (the default) or specify new metadata. If this header isn't specified, COPY is the default behavior.

General purpose bucket - For general purpose buckets, when you grant permissions, you can use the `s3:x-amz-metadata-directive` condition key to enforce certain metadata behavior when objects are uploaded. For more information, see [Amazon S3 condition key examples](#) in the *Amazon S3 User Guide*.

Note

`x-amz-website-redirect-location` is unique to each object and is not copied when using the `x-amz-metadata-directive` header. To copy the value, you must specify `x-amz-website-redirect-location` in the request header.

Valid Values: COPY | REPLACE

x-amz-object-annotation-directive

Specifies whether you want to copy annotations from the source object or exclude them. If this header isn't specified, COPY is the default behavior.

Valid Values: COPY | EXCLUDE

You can specify this directive as either an HTTP header (`x-amz-object-annotation-directive`) or as a query string parameter. Use the query string form when generating presigned URLs that need to control annotation copy behavior.

When set to COPY, you must have `s3:GetObjectAnnotation` permission on the source object and `s3:PutObjectAnnotation` permission on the destination. Each annotation copied is

billed as a separate PUT request. If annotations on the source are modified during the copy, Amazon S3 returns a retryable error.

 Note

For directory buckets, annotations are not supported. Use EXCLUDE to copy objects to directory buckets without errors. If you specify COPY for a directory bucket, the request returns HTTP 501 (Not Implemented).

 Note

When you copy objects using multipart upload (for example, when the AWS CLI or AWS SDKs use Transfer Manager for objects larger than approximately 8 MB), annotations are not copied by default. To include annotations, specify `--copy-props default` in the AWS CLI or the equivalent SDK configuration. With this opt-in, the SDK reads source annotations, completes the multipart upload, and then writes each annotation to the destination. Between the upload completion and the last annotation write, the destination object exists without all its annotations.

Valid Values: COPY | EXCLUDE

x-amz-object-lock-event-hold

The event hold status to apply to the object copy. Set to ON to enable or OFF to disable.

 Note

This functionality is not supported for directory buckets.

Valid Values: ON | OFF

x-amz-object-lock-event-hold-duration-days

The event hold duration in days to apply to the object copy. You cannot specify a duration in both days and years.

Note

This functionality is not supported for directory buckets.

x-amz-object-lock-event-hold-duration-years

The event hold duration in years to apply to the object copy. You cannot specify a duration in both days and years.

Note

This functionality is not supported for directory buckets.

x-amz-object-lock-legal-hold

Specifies whether you want to apply a legal hold to the object copy.

Note

This functionality is not supported for directory buckets.

Valid Values: ON | OFF

x-amz-object-lock-mode

The Object Lock mode that you want to apply to the object copy.

Note

This functionality is not supported for directory buckets.

Valid Values: GOVERNANCE | COMPLIANCE

x-amz-object-lock-retain-until-date

The date and time when you want the Object Lock of the object copy to expire.

Note

This functionality is not supported for directory buckets.

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

x-amz-server-side-encryption

The server-side encryption algorithm used when storing this object in Amazon S3. Unrecognized or unsupported values won't write a destination object and will receive a 400 Bad Request response.

Amazon S3 automatically encrypts all new objects that are copied to an S3 bucket. When copying an object, if you don't specify encryption information in your copy request, the encryption setting of the target object is set to the default encryption configuration of the destination bucket. By default, all buckets have a base level of encryption configuration that uses server-side encryption with Amazon S3 managed keys (SSE-S3). If the destination bucket has a different default encryption configuration, Amazon S3 uses the corresponding encryption key to encrypt the target object copy.

With server-side encryption, Amazon S3 encrypts your data as it writes your data to disks in its data centers and decrypts the data when you access it. For more information about server-side encryption, see [Using Server-Side Encryption](#) in the *Amazon S3 User Guide*.

General purpose buckets

- For general purpose buckets, there are the following supported options for server-side encryption: server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS), dual-layer server-side encryption with AWS KMS keys (DSSE-KMS), and server-side encryption with customer-provided encryption keys (SSE-C). Amazon S3 uses the corresponding KMS key, or a customer-provided key to encrypt the target object copy.
- When you perform a `CopyObject` operation, if you want to use a different type of encryption setting for the target object, you can specify appropriate encryption-related headers to encrypt the target object with an Amazon S3 managed key, a KMS key, or a customer-provided key. If the encryption setting in your request is different from the default encryption configuration of the destination bucket, the encryption setting in your request takes precedence.

Directory buckets

- For directory buckets, there are only two supported options for server-side encryption: server-side encryption with Amazon S3 managed keys (SSE-S3) (AES256) and server-side encryption with AWS KMS keys (SSE-KMS) (`aws:kms`). We recommend that the bucket's default encryption uses the desired encryption configuration and you don't override the bucket default encryption in your `CreateSession` requests or `PUT` object requests. Then, new objects are automatically encrypted with the desired encryption settings. For more information, see [Protecting data with server-side encryption](#) in the *Amazon S3 User Guide*. For more information about the encryption overriding behaviors in directory buckets, see [Specifying server-side encryption with AWS KMS for new object uploads](#).
- To encrypt new object copies to a directory bucket with SSE-KMS, we recommend you specify SSE-KMS as the directory bucket's default encryption configuration with a KMS key (specifically, a [customer managed key](#)). The [AWS managed key](#) (`aws/s3`) isn't supported. Your SSE-KMS configuration can only support 1 [customer managed key](#) per directory bucket for the lifetime of the bucket. After you specify a customer managed key for SSE-KMS, you can't override the customer managed key for the bucket's SSE-KMS configuration. Then, when you perform a `CopyObject` operation and want to specify server-side encryption settings for new object copies with SSE-KMS in the encryption-related request headers, you must ensure the encryption key is the same customer managed key that you specified for the directory bucket's default encryption configuration.
- **S3 access points for Amazon FSx** - When accessing data stored in Amazon FSx file systems using S3 access points, the only valid server side encryption option is `aws:fsx`. All Amazon FSx file systems have encryption configured by default and are encrypted at rest. Data is

automatically encrypted before being written to the file system, and automatically decrypted as it is read. These processes are handled transparently by Amazon FSx.

Valid Values: AES256 | aws:fsx | aws:backup | aws:kms | aws:kms:dsse

[x-amz-server-side-encryption-aws-kms-key-id](#)

Specifies the AWS KMS key ID (Key ID, Key ARN, or Key Alias) to use for object encryption. All GET and PUT requests for an object protected by AWS KMS will fail if they're not made via SSL or using SigV4. For information about configuring any of the officially supported AWS SDKs and AWS CLI, see [Specifying the Signature Version in Request Authentication](#) in the *Amazon S3 User Guide*.

Directory buckets - To encrypt data using SSE-KMS, it's recommended to specify the `x-amz-server-side-encryption` header to `aws:kms`. Then, the `x-amz-server-side-encryption-aws-kms-key-id` header implicitly uses the bucket's default KMS customer managed key ID. If you want to explicitly set the `x-amz-server-side-encryption-aws-kms-key-id` header, it must match the bucket's default customer managed key (using key ID or ARN, not alias). Your SSE-KMS configuration can only support 1 [customer managed key](#) per directory bucket's lifetime. The [AWS managed key](#) (`aws/s3`) isn't supported. Incorrect key specification results in an HTTP 400 Bad Request error.

[x-amz-server-side-encryption-bucket-key-enabled](#)

Specifies whether Amazon S3 should use an S3 Bucket Key for object encryption with server-side encryption using AWS Key Management Service (AWS KMS) keys (SSE-KMS). If a target object uses SSE-KMS, you can enable an S3 Bucket Key for the object.

Setting this header to `true` causes Amazon S3 to use an S3 Bucket Key for object encryption with SSE-KMS. Specifying this header with a COPY action doesn't affect bucket-level settings for S3 Bucket Key.

For more information, see [Amazon S3 Bucket Keys](#) in the *Amazon S3 User Guide*.

Note

Directory buckets - S3 Bucket Keys aren't supported, when you copy SSE-KMS encrypted objects from general purpose buckets to directory buckets, from directory buckets to general purpose buckets, or between directory buckets, through [CopyObject](#). In this case, Amazon S3 makes a call to AWS KMS every time a copy request is made for a KMS-encrypted object.

x-amz-server-side-encryption-context

Specifies the AWS KMS Encryption Context as an additional encryption context to use for the destination object encryption. The value of this header is a base64-encoded UTF-8 string holding JSON with the encryption context key-value pairs.

General purpose buckets - This value must be explicitly added to specify encryption context for CopyObject requests if you want an additional encryption context for your destination object. The additional encryption context of the source object won't be copied to the destination object. For more information, see [Encryption context](#) in the *Amazon S3 User Guide*.

Directory buckets - You can optionally provide an explicit encryption context value. The value must match the default encryption context - the bucket Amazon Resource Name (ARN). An additional encryption context value is not supported.

x-amz-server-side-encryption-customer-algorithm

Specifies the algorithm to use when encrypting the object (for example, AES256).

When you perform a CopyObject operation, if you want to use a different type of encryption setting for the target object, you can specify appropriate encryption-related headers to encrypt the target object with an Amazon S3 managed key, a KMS key, or a customer-provided key. If the encryption setting in your request is different from the default encryption configuration of the destination bucket, the encryption setting in your request takes precedence.

Note

This functionality is not supported when the destination bucket is a directory bucket.

x-amz-server-side-encryption-customer-key

Specifies the customer-provided encryption key for Amazon S3 to use in encrypting data. This value is used to store the object and then it is discarded. Amazon S3 does not store the encryption key. The key must be appropriate for use with the algorithm specified in the x-amz-server-side-encryption-customer-algorithm header.

Note

This functionality is not supported when the destination bucket is a directory bucket.

x-amz-server-side-encryption-customer-key-MD5

Specifies the 128-bit MD5 digest of the encryption key according to RFC 1321. Amazon S3 uses this header for a message integrity check to ensure that the encryption key was transmitted without error.

Note

This functionality is not supported when the destination bucket is a directory bucket.

x-amz-source-expected-bucket-owner

The account ID of the expected source bucket owner. If the account ID that you provide does not match the actual owner of the source bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-storage-class

If the `x-amz-storage-class` header is not used, the copied object will be stored in the STANDARD Storage Class by default. The STANDARD storage class provides high durability and high availability. Depending on performance needs, you can specify a different Storage Class.

Note

- **Directory buckets** - Directory buckets only support EXPRESS_ONEZONE (the S3 Express One Zone storage class) in Availability Zones and ONEZONE_IA (the S3 One Zone-Infrequent Access storage class) in Dedicated Local Zones. Unsupported storage class values won't write a destination object and will respond with the HTTP status code 400 Bad Request.
- **Amazon S3 on Outposts** - S3 on Outposts only uses the OUTPOSTS Storage Class.

You can use the `CopyObject` action to change the storage class of an object that is already stored in Amazon S3 by using the `x-amz-storage-class` header. For more information, see [Storage Classes](#) in the *Amazon S3 User Guide*.

Before using an object as a source object for the copy operation, you must restore a copy of it if it meets any of the following conditions:

- The storage class of the source object is GLACIER or DEEP_ARCHIVE.

- The storage class of the source object is `INTELLIGENT_TIERING` and it's [S3 Intelligent-Tiering access tier](#) is `Archive Access` or `Deep Archive Access`.

For more information, see [RestoreObject](#) and [Copying Objects](#) in the *Amazon S3 User Guide*.

Valid Values: `STANDARD` | `REDUCED_REDUNDANCY` | `STANDARD_IA` | `ONEZONE_IA` | `INTELLIGENT_TIERING` | `GLACIER` | `DEEP_ARCHIVE` | `OUTPOSTS` | `GLACIER_IR` | `SNOW` | `EXPRESS_ONEZONE` | `FSX_OPENZFS` | `FSX_ONTAP` | `AWS_BACKUP_WARM` | `AWS_BACKUP_LOW_COST_WARM`

[x-amz-tagging](#)

The tag-set for the object copy in the destination bucket. This value must be used in conjunction with the `x-amz-tagging-directive` if you choose `REPLACE` for the `x-amz-tagging-directive`. If you choose `COPY` for the `x-amz-tagging-directive`, you don't need to set the `x-amz-tagging` header, because the tag-set will be copied from the source object directly. The tag-set must be encoded as URL Query parameters.

The default value is the empty value.

Note

Directory buckets - For directory buckets in a `CopyObject` operation, only the empty tag-set is supported. Any requests that attempt to write non-empty tags into directory buckets will receive a `501 Not Implemented` status code. When the destination bucket is a directory bucket, you will receive a `501 Not Implemented` response in any of the following situations:

- When you attempt to `COPY` the tag-set from an S3 source object that has non-empty tags.
- When you attempt to `REPLACE` the tag-set of a source object and set a non-empty value to `x-amz-tagging`.
- When you don't set the `x-amz-tagging-directive` header and the source object has non-empty tags. This is because the default value of `x-amz-tagging-directive` is `COPY`.

Because only the empty tag-set is supported for directory buckets in a `CopyObject` operation, the following situations are allowed:

- When you attempt to `COPY` the tag-set from a directory bucket source object that has no tags to a general purpose bucket. It copies an empty tag-set to the destination object.

- When you attempt to REPLACE the tag-set of a directory bucket source object and set the `x-amz-tagging` value of the directory bucket destination object to empty.
- When you attempt to REPLACE the tag-set of a general purpose bucket source object that has non-empty tags and set the `x-amz-tagging` value of the directory bucket destination object to empty.
- When you attempt to REPLACE the tag-set of a directory bucket source object and don't set the `x-amz-tagging` value of the directory bucket destination object. This is because the default value of `x-amz-tagging` is the empty value.

`x-amz-tagging-directive`

Specifies whether the object tag-set is copied from the source object or replaced with the tag-set that's provided in the request.

The default value is `COPY`.

Note

Directory buckets - For directory buckets in a `CopyObject` operation, only the empty tag-set is supported. Any requests that attempt to write non-empty tags into directory buckets will receive a `501 Not Implemented` status code. When the destination bucket is a directory bucket, you will receive a `501 Not Implemented` response in any of the following situations:

- When you attempt to `COPY` the tag-set from an S3 source object that has non-empty tags.
- When you attempt to `REPLACE` the tag-set of a source object and set a non-empty value to `x-amz-tagging`.
- When you don't set the `x-amz-tagging-directive` header and the source object has non-empty tags. This is because the default value of `x-amz-tagging-directive` is `COPY`.

Because only the empty tag-set is supported for directory buckets in a `CopyObject` operation, the following situations are allowed:

- When you attempt to `COPY` the tag-set from a directory bucket source object that has no tags to a general purpose bucket. It copies an empty tag-set to the destination object.

- When you attempt to REPLACE the tag-set of a directory bucket source object and set the `x-amz-tagging` value of the directory bucket destination object to empty.
- When you attempt to REPLACE the tag-set of a general purpose bucket source object that has non-empty tags and set the `x-amz-tagging` value of the directory bucket destination object to empty.
- When you attempt to REPLACE the tag-set of a directory bucket source object and don't set the `x-amz-tagging` value of the directory bucket destination object. This is because the default value of `x-amz-tagging` is the empty value.

Valid Values: COPY | REPLACE

`x-amz-website-redirect-location`

If the destination bucket is configured as a website, redirects requests for this object copy to another object in the same bucket or to an external URL. Amazon S3 stores the value of this header in the object metadata. This value is unique to each object and is not copied when using the `x-amz-metadata-directive` header. Instead, you may opt to provide this header in combination with the `x-amz-metadata-directive` header.

Note

This functionality is not supported for directory buckets.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
x-amz-expiration: Expiration
x-amz-copy-source-version-id: CopySourceVersionId
x-amz-version-id: VersionId
x-amz-server-side-encryption: ServerSideEncryption
x-amz-server-side-encryption-customer-algorithm: SSECustomerAlgorithm
x-amz-server-side-encryption-customer-key-MD5: SSECustomerKeyMD5
x-amz-server-side-encryption-aws-kms-key-id: SSEKMSKeyId
x-amz-server-side-encryption-context: SSEKMSEncryptionContext
```



```

x-amz-server-side-encryption-bucket-key-enabled: BucketKeyEnabled
x-amz-request-charged: RequestCharged
<?xml version="1.0" encoding="UTF-8"?>
<CopyObjectResult>
  <ETag>string</ETag>
  <LastModified>timestamp</LastModified>
  <ChecksumType>string</ChecksumType>
  <ChecksumCRC32>string</ChecksumCRC32>
  <ChecksumCRC32C>string</ChecksumCRC32C>
  <ChecksumCRC64NVME>string</ChecksumCRC64NVME>
  <ChecksumSHA1>string</ChecksumSHA1>
  <ChecksumSHA256>string</ChecksumSHA256>
  <ChecksumSHA512>string</ChecksumSHA512>
  <ChecksumMD5>string</ChecksumMD5>
  <ChecksumXXHASH64>string</ChecksumXXHASH64>
  <ChecksumXXHASH3>string</ChecksumXXHASH3>
  <ChecksumXXHASH128>string</ChecksumXXHASH128>
</CopyObjectResult>

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

x-amz-copy-source-version-id

Version ID of the source object that was copied.

Note

This functionality is not supported when the source object is in a directory bucket.

x-amz-expiration

If the object expiration is configured, the response includes this header.

Note

Object expiration information is not returned in directory buckets and this header returns the value "NotImplemented" in all responses for directory buckets.

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

x-amz-server-side-encryption

The server-side encryption algorithm used when you store this object in Amazon S3 or Amazon FSx.

Note

When accessing data stored in Amazon FSx file systems using S3 access points, the only valid server side encryption option is `aws:fsx`.

Valid Values: AES256 | `aws:fsx` | `aws:backup` | `aws:kms` | `aws:kms:dsse`

x-amz-server-side-encryption-aws-kms-key-id

If present, indicates the ID of the KMS key that was used for object encryption.

x-amz-server-side-encryption-bucket-key-enabled

Indicates whether the copied object uses an S3 Bucket Key for server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS).

x-amz-server-side-encryption-context

If present, indicates the AWS KMS Encryption Context to use for object encryption. The value of this header is a Base64 encoded UTF-8 string holding JSON with the encryption context key-value pairs.

x-amz-server-side-encryption-customer-algorithm

If server-side encryption with a customer-provided encryption key was requested, the response will include this header to confirm the encryption algorithm that's used.

Note

This functionality is not supported for directory buckets.

x-amz-server-side-encryption-customer-key-MD5

If server-side encryption with a customer-provided encryption key was requested, the response will include this header to provide the round-trip message integrity verification of the customer-provided encryption key.

Note

This functionality is not supported for directory buckets.

x-amz-version-id

Version ID of the newly created copy.

Note

This functionality is not supported for directory buckets.

The following data is returned in XML format by the service.

CopyObjectResult

Root level tag for the CopyObjectResult parameters.

Required: Yes

ChecksumCRC32

The Base64 encoded, 32-bit CRC32 checksum of the object. This checksum is only present if the checksum was uploaded with the object. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

ChecksumCRC32C

The Base64 encoded, 32-bit CRC32C checksum of the object. This checksum is only present if the checksum was uploaded with the object. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

ChecksumCRC64NVME

The Base64 encoded, 64-bit CRC64NVME checksum of the object. This checksum is present if the object being copied was uploaded with the CRC64NVME checksum algorithm, or if the object was uploaded without a checksum (and Amazon S3 added the default checksum, CRC64NVME, to the uploaded object). For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

ChecksumMD5

The Base64 encoded, 128-bit MD5 digest of the object. This checksum is only present if the object was uploaded with the MD5 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

ChecksumSHA1

The Base64 encoded, 160-bit SHA1 digest of the object. This checksum is only present if the checksum was uploaded with the object. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

ChecksumSHA256

The Base64 encoded, 256-bit SHA256 digest of the object. This checksum is only present if the checksum was uploaded with the object. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

ChecksumSHA512

The Base64 encoded, 512-bit SHA512 digest of the object. This checksum is only present if the object was uploaded with the SHA512 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

ChecksumType

The checksum type that is used to calculate the object's checksum value. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Valid Values: COMPOSITE | FULL_OBJECT

ChecksumXXHASH128

The Base64 encoded, 128-bit XXHASH128 checksum of the object. This checksum is only present if the object was uploaded with the XXHASH128 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

ChecksumXXHASH3

The Base64 encoded, 64-bit XXHASH3 checksum of the object. This checksum is only present if the object was uploaded with the XXHASH3 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

ChecksumXXHASH64

The Base64 encoded, 64-bit XXHASH64 checksum of the object. This checksum is only present if the object was uploaded with the XXHASH64 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

ETag

Returns the ETag of the new object. The ETag reflects only changes to the contents of an object, not its metadata.

Type: String

LastModified

Creation date of the object.

Type: Timestamp

Errors

ObjectNotInActiveTierError

The source object of the COPY action is not in the active tier and is only stored in Amazon S3 Glacier.

HTTP Status Code: 403

Examples

Sample Request for general purpose buckets

This example copies `my-image.jpg` into the `amzn-s3-demo-bucket` bucket, with the key name `my-second-image.jpg`.

```
PUT /my-second-image.jpg HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Wed, 28 Oct 2009 22:32:00 GMT
x-amz-copy-source: /amzn-s3-demo-bucket/my-image.jpg
Authorization: authorization string
```

Sample Response for general purpose buckets

This example illustrates one usage of `CopyObject`.

```

HTTP/1.1 200 OK
x-amz-id-2:
eftixk72aD6Ap51TnqcoF8eFidJG9Z/2mkiDFu8yU9AS1ed40pIszj7UDNEHGran
x-amz-request-id: 318BC8BC148832E5
x-amz-copy-source-version-id: 3/L4kqtJlcpXroDTDmJ+rmSpXd3dIbrHY
+MTRCxf3vjVBH40Nr8X8gdRQBpUMLUo
x-amz-version-id: QUpfdndhfd8438MNFdN93jdnJFkdmqnh893
Date: Wed, 28 Oct 2009 22:32:00 GMT
Connection: close
Server: AmazonS3

<CopyObjectResult>
  <LastModified>2009-10-12T17:50:30.000Z</LastModified>
  <ETag>"9b2cf535f27731c974343645a3985328"</ETag>
</CopyObjectResult>

```

Sample Request for general purpose buckets: Copying a specified version of an object

The following request copies the `my-image.jpg` key with the specified version ID, copies it into the `amzn-s3-demo-bucket` bucket, and gives it the `my-second-image.jpg` key.

```

PUT /my-second-image.jpg HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Wed, 28 Oct 2009 22:32:00 GMT
x-amz-copy-source: /amzn-s3-demo-bucket/my-image.jpg?versionId=3/
L4kqtJlcpXroDTDmJ+rmSpXd3dIbrHY+MTRCxf3vjVBH40Nr8X8gdRQBpUMLUo
Authorization: authorization string

```

Success Response for general purpose buckets: Copying a versioned object into a version-enabled bucket

The following response shows that an object was copied into a target bucket where versioning is enabled.

```

HTTP/1.1 200 OK
x-amz-id-2:
eftixk72aD6Ap51TnqcoF8eFidJG9Z/2mkiDFu8yU9AS1ed40pIszj7UDNEHGran

```

```

x-amz-request-id: 318BC8BC148832E5
x-amz-version-id: QUpfdndhfd8438MNFDN93jdnJFkdmqnh893
x-amz-copy-source-version-id: 09df8234529fjs0dfi0w52935029wefdj
Date: Wed, 28 Oct 2009 22:32:00 GMT
Connection: close
Server: AmazonS3

<?xml version="1.0" encoding="UTF-8"?>
<CopyObjectResult>
  <LastModified>2009-10-12T17:50:30.000Z</LastModified>
  <ETag>"9b2cf535f27731c974343645a3985328"</ETag>
</CopyObjectResult>

```

Success Response for general purpose buckets: Copying a versioned object into a version-suspended bucket

The following response shows that an object was copied into a target bucket where versioning is suspended. The parameter `VersionId` does not appear.

```

HTTP/1.1 200 OK
x-amz-id-2:
eftixk72aD6Ap51TnqcoF8eFidJG9Z/2mkiDFu8yU9AS1ed40pIszj7UDNEHGran
x-amz-request-id: 318BC8BC148832E5
x-amz-copy-source-version-id: 3/L4kqtJlcpXroDTDmJ+rmSpXd3dIbrHY
+MTRCxf3vjVBH40Nr8X8gdRQBpUMLUo
Date: Wed, 28 Oct 2009 22:32:00 GMT
Connection: close
Server: AmazonS3

<?xml version="1.0" encoding="UTF-8"?>
<CopyObjectResult>
  <LastModified>2009-10-28T22:32:00</LastModified>
  <ETag>"9b2cf535f27731c974343645a3985328"</ETag>
</CopyObjectResult>

```

Sample Request for general purpose buckets: Copy from unencrypted object to an object encrypted with server-side encryption with customer-provided encryption keys

The following example specifies the HTTP PUT header to copy an unencrypted object to an object encrypted with server-side encryption with customer-provided encryption keys (SSE-C).


```
PUT /exampleDestinationObject HTTP/1.1
Host: amzn-s3-demo-destination-bucket.s3.<Region>.amazonaws.com
x-amz-server-side-encryption-customer-algorithm: AES256
x-amz-server-side-encryption-customer-key: Base64(YourKey)
x-amz-server-side-encryption-customer-key-MD5 : Base64(MD5(YourKey))
x-amz-metadata-directive: metadata_directive
x-amz-copy-source: /amzn-s3-demo-source-bucket/exampleSourceObject
x-amz-copy-source-if-match: etag
x-amz-copy-source-if-none-match: etag
x-amz-copy-source-if-unmodified-since: time_stamp
x-amz-copy-source-if-modified-since: time_stamp
```

```
<request metadata>
```

```
Authorization: authorization string (see Authenticating Requests (AWS
Signature Version 4))
```

```
Date: date
```

Sample Request for general purpose buckets: Copy from an object encrypted with SSE-C to an object encrypted with SSE-C

The following example specifies the HTTP PUT header to copy an object encrypted with server-side encryption with customer-provided encryption keys to an object encrypted with server-side encryption with customer-provided encryption keys for key rotation.

Note

If you have server-side encryption with customer-provided keys (SSE-C) blocked for your general purpose bucket, you will get an HTTP 403 Access Denied error when you specify the SSE-C request headers while writing new data to your bucket. For more information, see [Blocking or unblocking SSE-C for a general purpose bucket](#).

```
PUT /exampleDestinationObject HTTP/1.1
Host: amzn-s3-demo-destination-bucket.s3.<Region>.amazonaws.com
x-amz-server-side-encryption-customer-algorithm: AES256
x-amz-server-side-encryption-customer-key: Base64(NewKey)
x-amz-server-side-encryption-customer-key-MD5: Base64(MD5(NewKey))
x-amz-metadata-directive: metadata_directive
```

```
x-amz-copy-source: /amzn-s3-demo-source-bucket/sourceObject
x-amz-copy-source-if-match: etag
x-amz-copy-source-if-none-match: etag
x-amz-copy-source-if-unmodified-since: time_stamp
x-amz-copy-source-if-modified-since: time_stamp
x-amz-copy-source-server-side-encryption-customer-algorithm: AES256
x-amz-copy-source-server-side-encryption-customer-key: Base64(OldKey)
x-amz-copy-source-server-side-encryption-customer-key-MD5:
Base64(MD5(OldKey))

<request metadata>
  Authorization: authorization string (see Authenticating Requests (AWS
Signature Version 4))
  Date: date
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

CreateBucket

Service: Amazon S3

Note

This action creates an Amazon S3 bucket. To create an Amazon S3 on Outposts bucket, see [CreateBucket](#).

Creates a new S3 bucket. To create a bucket, you must set up Amazon S3 and have a valid AWS Access Key ID to authenticate requests. Anonymous requests are never allowed to create buckets. By creating the bucket, you become the bucket owner.

There are two types of buckets: general purpose buckets and directory buckets. For more information about these bucket types, see [Creating, configuring, and working with Amazon S3 buckets](#) in the *Amazon S3 User Guide*.

General purpose buckets exist in a global namespace, which means that each bucket name must be unique across all AWS accounts in all the AWS Regions within a partition. A partition is a grouping of Regions. AWS currently has four partitions: aws (Standard Regions), aws-cn (China Regions), aws-us-gov (AWS GovCloud (US)), and aws-eusc (European Sovereign Cloud). When you create a general purpose bucket, you can choose to create a bucket in the shared global namespace or you can choose to create a bucket in your account regional namespace. Your account regional namespace is a subdivision of the global namespace that only your account can create buckets in. For more information on account regional namespaces, see [Namespaces for general purpose buckets](#).

Note

- **General purpose buckets** - If you send your CreateBucket request to the `s3.amazonaws.com` global endpoint, the request goes to the us-east-1 Region. So the signature calculations in Signature Version 4 must use us-east-1 as the Region, even if the location constraint in the request specifies another Region where the bucket is to be created. If you create a bucket in a Region other than US East (N. Virginia), your application must be able to handle 307 redirect. For more information, see [Virtual hosting of buckets](#) in the *Amazon S3 User Guide*.

- **Directory buckets** - For directory buckets, you must make requests for this API operation to the Regional endpoint. These endpoints support path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name` . Virtual-hosted-style requests aren't supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

Permissions

- **General purpose bucket permissions** - In addition to the `s3:CreateBucket` permission, the following permissions are required in a policy when your `CreateBucket` request includes specific headers:
 - **Access control lists (ACLs)** - In your `CreateBucket` request, if you specify an access control list (ACL) and set it to `public-read`, `public-read-write`, `authenticated-read`, or if you explicitly specify any other custom ACLs, both `s3:CreateBucket` and `s3:PutBucketAcl` permissions are required. In your `CreateBucket` request, if you set the ACL to `private`, or if you don't specify any ACLs, only the `s3:CreateBucket` permission is required.
 - **Object Lock** - In your `CreateBucket` request, if you set `x-amz-bucket-object-lock-enabled` to `true`, the `s3:PutBucketObjectLockConfiguration` and `s3:PutBucketVersioning` permissions are required.
 - **S3 Object Ownership** - If your `CreateBucket` request includes the `x-amz-object-ownership` header, then the `s3:PutBucketOwnershipControls` permission is required.

Important

To set an ACL on a bucket as part of a `CreateBucket` request, you must explicitly set S3 Object Ownership for the bucket to a different value than the default, `BucketOwnerEnforced`. Additionally, if your desired bucket ACL grants public access, you must first create the bucket (without the bucket ACL) and then explicitly disable Block Public Access on the bucket before using `PutBucketAcl` to set the ACL. If you try to create a bucket with a public ACL, the request will fail. For the majority of modern use cases in S3, we recommend that you keep all Block Public Access settings enabled and keep ACLs disabled. If you would like to share data with users outside of your account, you can use bucket policies as needed. For

more information, see [Controlling ownership of objects and disabling ACLs for your bucket](#) and [Blocking public access to your Amazon S3 storage](#) in the *Amazon S3 User Guide*.

- **S3 Block Public Access** - If your specific use case requires granting public access to your S3 resources, you can disable Block Public Access. Specifically, you can create a new bucket with Block Public Access enabled, then separately call the [DeletePublicAccessBlock](#) API. To use this operation, you must have the `s3:PutBucketPublicAccessBlock` permission. For more information about S3 Block Public Access, see [Blocking public access to your Amazon S3 storage](#) in the *Amazon S3 User Guide*.
- **Directory bucket permissions** - You must have the `s3express:CreateBucket` permission in an IAM identity-based policy instead of a bucket policy. Cross-account access to this API operation isn't supported. This operation can only be performed by the AWS account that owns the resource. For more information about directory bucket policies and permissions, see [AWS Identity and Access Management \(IAM\) for S3 Express One Zone](#) in the *Amazon S3 User Guide*.

 Important

The permissions for ACLs, Object Lock, S3 Object Ownership, and S3 Block Public Access are not supported for directory buckets. For directory buckets, all Block Public Access settings are enabled at the bucket level and S3 Object Ownership is set to Bucket owner enforced (ACLs disabled). These settings can't be modified.

For more information about permissions for creating and working with directory buckets, see [Directory buckets](#) in the *Amazon S3 User Guide*. For more information about supported S3 features for directory buckets, see [Features of S3 Express One Zone](#) in the *Amazon S3 User Guide*.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `s3express-control.region-code.amazonaws.com`.

The following operations are related to `CreateBucket`:

- [PutObject](#)
- [DeleteBucket](#)

⚠ Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT / HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-acl: ACL
x-amz-grant-full-control: GrantFullControl
x-amz-grant-read: GrantRead
x-amz-grant-read-acp: GrantReadACP
x-amz-grant-write: GrantWrite
x-amz-grant-write-acp: GrantWriteACP
x-amz-bucket-object-lock-enabled: ObjectLockEnabledForBucket
x-amz-object-ownership: ObjectOwnership
x-amz-bucket-namespace: BucketNamespace
<?xml version="1.0" encoding="UTF-8"?>
<CreateBucketConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <LocationConstraint>string</LocationConstraint>
  <Location>
    <Name>string</Name>
    <Type>string</Type>
  </Location>
  <Bucket>
    <DataRedundancy>string</DataRedundancy>
    <Type>string</Type>
  </Bucket>
  <Tags>
    <Tag>
      <Key>string</Key>
      <Value>string</Value>
    </Tag>
  </Tags>
</CreateBucketConfiguration>
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket to create.

General purpose buckets - For information about bucket naming restrictions, see [Bucket naming rules](#) in the *Amazon S3 User Guide*.

Directory buckets - When you use this operation with a directory bucket, you must use path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name`. Virtual-hosted-style requests aren't supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must also follow the format `bucket-base-name--zone-id--x-s3` (for example, `DOC-EXAMPLE-BUCKET--usw2-az1--x-s3`). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*.

Required: Yes

x-amz-acl

The canned ACL to apply to the bucket.

Note

This functionality is not supported for directory buckets.

Valid Values: `private` | `public-read` | `public-read-write` | `authenticated-read`

x-amz-bucket-namespace

Specifies the namespace where you want to create your general purpose bucket. When you create a general purpose bucket, you can choose to create a bucket in the shared global namespace or you can choose to create a bucket in your account regional namespace. Your account regional namespace is a subdivision of the global namespace that only your account can create buckets in. For more information on bucket namespaces, see [Namespaces for general purpose buckets](#).

General purpose buckets in your account regional namespace must follow a specific naming convention. These buckets consist of a bucket name prefix that you create, and a suffix that contains your 12-digit AWS Account ID, the AWS Region code, and ends with `-an`. Bucket names must follow the format `bucket-name-prefix-accountId-region-an` (for example, `amzn-`

s3-demo-bucket-111122223333-us-west-2-an). For information about bucket naming restrictions, see [Account regional namespace naming rules](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: account-regional | global

x-amz-bucket-object-lock-enabled

Specifies whether you want S3 Object Lock to be enabled for the new bucket.

Note

This functionality is not supported for directory buckets.

x-amz-grant-full-control

Allows grantee the read, write, read ACP, and write ACP permissions on the bucket.

Note

This functionality is not supported for directory buckets.

x-amz-grant-read

Allows grantee to list the objects in the bucket.

Note

This functionality is not supported for directory buckets.

x-amz-grant-read-acp

Allows grantee to read the bucket ACL.

Note

This functionality is not supported for directory buckets.

x-amz-grant-write

Allows grantee to create new objects in the bucket.

For the bucket and object owners of existing objects, also allows deletions and overwrites of those objects.

Note

This functionality is not supported for directory buckets.

x-amz-grant-write-acp

Allows grantee to write the ACL for the applicable bucket.

Note

This functionality is not supported for directory buckets.

x-amz-object-ownership

The container element for object ownership for a bucket's ownership controls.

`BucketOwnerPreferred` - Objects uploaded to the bucket change ownership to the bucket owner if the objects are uploaded with the `bucket-owner-full-control` canned ACL.

`ObjectWriter` - The uploading account will own the object if the object is uploaded with the `bucket-owner-full-control` canned ACL.

`BucketOwnerEnforced` - Access control lists (ACLs) are disabled and no longer affect permissions. The bucket owner automatically owns and has full control over every object in the bucket. The bucket only accepts PUT requests that don't specify an ACL or specify bucket owner full control ACLs (such as the predefined `bucket-owner-full-control` canned ACL or a custom ACL in XML format that grants the same permissions).

By default, `ObjectOwnership` is set to `BucketOwnerEnforced` and ACLs are disabled. We recommend keeping ACLs disabled, except in uncommon use cases where you must control access for each object individually. For more information about S3 Object Ownership, see [Controlling ownership of objects and disabling ACLs for your bucket](#) in the *Amazon S3 User Guide*.

 Note

This functionality is not supported for directory buckets. Directory buckets use the bucket owner enforced setting for S3 Object Ownership.

Valid Values: `BucketOwnerPreferred` | `ObjectWriter` | `BucketOwnerEnforced`

Request Body

The request accepts the following data in XML format.

CreateBucketConfiguration

Root level tag for the `CreateBucketConfiguration` parameters.

Required: Yes

Bucket

Specifies the information about the bucket that will be created.

 Note

This functionality is only supported by directory buckets.

Type: [BucketInfo](#) data type

Required: No

Location

Specifies the location where the bucket will be created.

Directory buckets - The location type is Availability Zone or Local Zone. To use the Local Zone location type, your account must be enabled for Local Zones. Otherwise, you get an HTTP 403 Forbidden error with the error code `AccessDenied`. To learn more, see [Enable accounts for Local Zones](#) in the *Amazon S3 User Guide*.

 Note

This functionality is only supported by directory buckets.

Type: [LocationInfo](#) data type

Required: No

[LocationConstraint](#)

Specifies the Region where the bucket will be created. You might choose a Region to optimize latency, minimize costs, or address regulatory requirements. For example, if you reside in Europe, you will probably find it advantageous to create buckets in the Europe (Ireland) Region.

If you don't specify a Region, the bucket is created in the US East (N. Virginia) Region (`us-east-1`) by default. Configurations using the value `EU` will create a bucket in `eu-west-1`.

For a list of the valid values for all of the AWS Regions, see [Regions and Endpoints](#).

 Note

This functionality is not supported for directory buckets.

Type: String

Valid Values: `af-south-1` | `ap-east-1` | `ap-east-2` | `ap-northeast-1` | `ap-northeast-2` | `ap-northeast-3` | `ap-south-1` | `ap-south-2` | `ap-southeast-1` | `ap-southeast-2` | `ap-southeast-3` | `ap-southeast-4` | `ap-southeast-5` | `ap-southeast-6` | `ap-southeast-7` | `ca-central-1` | `ca-west-1` | `cn-north-1` | `cn-northwest-1` | `EU` | `eu-central-1` | `eu-central-2` | `eu-north-1` | `eu-south-1` | `eu-south-2` | `eu-west-1` | `eu-west-2` | `eu-west-3` | `il-central-1` | `me-central-1` | `me-south-1` | `mx-central-1` | `sa-east-1` | `us-east-2` | `us-gov-east-1` | `us-gov-west-1` | `us-west-1` | `us-west-2`

Required: No

Tags

An array of tags that you can apply to the bucket that you're creating. Tags are key-value pairs of metadata used to categorize and organize your buckets, track costs, and control access.

You must have the `s3:TagResource` permission to create a general purpose bucket with tags or the `s3express:TagResource` permission to create a directory bucket with tags.

When creating buckets with tags, note that tag-based conditions using `aws:ResourceTag` and `s3:BucketTag` condition keys are applicable only after ABAC is enabled on the bucket. To learn more, see [Enabling ABAC in general purpose buckets](#).

Type: Array of [Tag](#) data types

Required: No

Response Syntax

```
HTTP/1.1 200
Location: Location
x-amz-bucket-arn: BucketArn
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

[Location](#)

A forward slash followed by the name of the bucket for all account regional namespace buckets and all global general purpose buckets created in us-east-1. For example, `/amzn-s3-demo-bucket`. For global general purpose buckets created in other AWS Regions, the Location field is the global endpoint URL. For example, `http://amzn-s3-demo-bucket.s3.amazonaws.com/`.

[x-amz-bucket-arn](#)

The Amazon Resource Name (ARN) of the S3 bucket. ARNs uniquely identify AWS resources across all of AWS.

Note

This parameter is only supported for S3 directory buckets. For more information, see [Using tags with directory buckets](#).

Length Constraints: Minimum length of 1. Maximum length of 128.

Pattern: `arn:[^:]+:(s3|s3express):.*`

Errors

BucketAlreadyExists

The requested bucket name is not available. The bucket namespace is shared by all users of the system. Select a different name and try again.

HTTP Status Code: 409

BucketAlreadyOwnedByYou

The bucket you tried to create already exists, and you own it. Amazon S3 returns this error in all AWS Regions except in the North Virginia Region. For legacy compatibility, if you re-create an existing bucket that you already own in the North Virginia Region, Amazon S3 returns 200 OK and resets the bucket access control lists (ACLs).

HTTP Status Code: 409

Examples

Sample Request for general purpose buckets

This request creates a bucket named `amzn-s3-demo-bucket`.

```
PUT / HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Content-Length: 0
Date: Wed, 01 Mar 2006 12:00:00 GMT
Authorization: authorization string
```

Sample Response for general purpose buckets

This example illustrates one usage of CreateBucket.

```
HTTP/1.1 200 OK
x-amz-id-2:
YgIPIfBiKa2bj0KMg95r/0zo3emzU4dzsD4rcKCHQUAdQkf3ShJT00pXUueF6QKo
x-amz-request-id: 236A8905248E5A01
Date: Wed, 01 Mar 2006 12:00:00 GMT

Location: /amzn-s3-demo-bucket
Content-Length: 0
Connection: close
Server: AmazonS3
```

Sample Request for general purpose buckets: Setting the Region of a bucket

The following request sets the Region for the bucket to Europe.

```
PUT / HTTP/1.1
Host: amzn-s3-demo-bucket.s3.amazonaws.com
Date: Wed, 12 Oct 2009 17:50:00 GMT
Authorization: authorization string
Content-Type: text/plain
Content-Length: 124

<CreateBucketConfiguration xmlns="http://s3.amazonaws.com/
doc/2006-03-01/">
  <LocationConstraint>EU</LocationConstraint>
</CreateBucketConfiguration >
```

Sample Request for general purpose buckets: Creating a bucket and applying the ObjectWriter setting for S3 Object Ownership.

This request creates a bucket and applies the ObjectWriter setting for Object Ownership.

```
PUT / HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
```

```
Content-Length: 0
x-amz-object-ownership: ObjectWriter
Date: Tue, 30 Nov 2021 12:00:00 GMT
Authorization: authorization string
```

Sample Response for general purpose buckets

This example illustrates one usage of CreateBucket.

```
HTTP/1.1 200 OK
x-amz-id-2:
YgIPIfBiKa2bj0KMg95r/0zo3emzU4dzsD4rcKCHQUAdQkf3ShJT00pXUueF6QKo
x-amz-request-id: 236A8905248E5A01
Date: Tue, 30 Nov 2021 12:00:00 GMT

Location: /amzn-s3-demo-bucket
Content-Length: 0
Connection: close
Server: AmazonS3
```

Sample Request for general purpose buckets: Creating a bucket and configuring access permissions explicitly

This request creates a bucket named `amzn-s3-demo-bucket` and grants WRITE permission to the AWS account identified by canonical id.

```
PUT HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
x-amz-date: Sat, 07 Apr 2012 00:54:40 GMT
Authorization: authorization string
x-amz-grant-write:
id="79a59df900b949e55d96a1e698fbacedfd6e09d98eacf8f8d5218e7cd47ef2be",
id="52b113e7a2f25102679df27bb0746532b8776c37c9fc7ca4ac94015369746b9f"
```

Sample Response for general purpose buckets

This example illustrates one usage of CreateBucket.

```
HTTP/1.1 200 OK
```

Sample Request for general purpose buckets: Creating a bucket and configuring access permission using a canned ACL

This request creates a bucket named `amzn-s3-demo-bucket` and sets the ACL to private.

```
PUT / HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Content-Length: 0
x-amz-acl: private
Date: Wed, 01 Mar 2006 12:00:00 GMT
Authorization: authorization string
```

Sample Response for general purpose buckets

This example illustrates one usage of `CreateBucket`.

```
HTTP/1.1 200 OK
x-amz-id-2:
YgIPIfBiKa2bj0KMg95r/0zo3emzU4dzsD4rcKCHQUAdQkf3ShJT00pXUueF6QKo
x-amz-request-id: 236A8905248E5A01
Date: Wed, 01 Mar 2006 12:00:00 GMT

Location: /amzn-s3-demo-bucket
Content-Length: 0
Connection: close
Server: AmazonS3
```

Sample Request for general purpose buckets: Creating a bucket in your account regional namespace

This request creates the general purpose bucket `amzn-s3-demo-bucket-012345678910-us-west-2-an` in the account regional namespace for AWS Account 012345678910 in the US West (Oregon) Region (`us-west-2`) Region.


```
PUT / HTTP/1.1
Host: amzn-s3-demo-bucket-012345678910-us-west-2-an.s3.us-
west-2.amazonaws.com
Content-Length: 0
x-amz-bucket-namespace: account-regional
Date: Wed, 01 Mar 2006 12:00:00 GMT
Authorization: authorization string
<CreateBucketConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <LocationConstraint>us-west-2</LocationConstraint>
</CreateBucketConfiguration >
```

Sample Response for general purpose buckets

This example illustrates one usage of CreateBucket.

```
HTTP/1.1 200 OK
x-amz-id-2:
YgIPIfBiKa2bj0KMg95r/0zo3emzU4dzsD4rcKCHQUAdQkf3ShJT00pXUueF6QKo
x-amz-request-id: 236A8905248E5A01
Date: Wed, 01 Mar 2006 12:00:00 GMT

Location: /amzn-s3-demo-bucket-012345678910-us-west-2-an
Content-Length: 0
Connection: close
Server: AmazonS3
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)

- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

CreateBucketMetadataConfiguration

Service: Amazon S3

Creates an S3 Metadata V2 metadata configuration for a general purpose bucket. For more information, see [Accelerating data discovery with S3 Metadata](#) in the *Amazon S3 User Guide*.

Permissions

To use this operation, you must have the following permissions. For more information, see [Setting up permissions for configuring metadata tables](#) in the *Amazon S3 User Guide*.

If you want to encrypt your metadata tables with server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS), you need additional permissions in your KMS key policy. For more information, see [Setting up permissions for configuring metadata tables](#) in the *Amazon S3 User Guide*.

If you also want to integrate your table bucket with AWS analytics services so that you can query your metadata table, you need additional permissions. For more information, see [Integrating Amazon S3 Tables with AWS analytics services](#) in the *Amazon S3 User Guide*.

To query your metadata tables, you need additional permissions. For more information, see [Permissions for querying metadata tables](#) in the *Amazon S3 User Guide*.

- `s3:CreateBucketMetadataTableConfiguration`

Note

The IAM policy action name is the same for the V1 and V2 API operations.

- `s3tables:CreateTableBucket`
- `s3tables:CreateNamespace`
- `s3tables:GetTable`
- `s3tables:CreateTable`
- `s3tables:PutTablePolicy`
- `s3tables:PutTableBucketPolicy`
- `s3tables:PutTableEncryption`
- `kms:DescribeKey`

- `iam:PassRole` - required if you include an `AnnotationTableConfiguration` with an IAM role.

The following operations are related to `CreateBucketMetadataConfiguration`:

- [DeleteBucketMetadataConfiguration](#)
- [GetBucketMetadataConfiguration](#)
- [UpdateBucketMetadataInventoryTableConfiguration](#)
- [UpdateBucketMetadataJournalTableConfiguration](#)
- [UpdateBucketMetadataAnnotationTableConfiguration](#)

If you include an `AnnotationTableConfiguration` with an IAM role, the role must have a trust policy that allows the Amazon S3 metadata service to assume it, and a permissions policy that grants the actions needed to read annotations from your bucket. The following examples show a trust policy and a permissions policy that you can adapt for your bucket and account.

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
POST /?metadataConfiguration HTTP/1.1
Host: Bucket.s3.amazonaws.com
Content-MD5: ContentMD5
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<MetadataConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <JournalTableConfiguration>
    <EncryptionConfiguration>
      <KmsKeyArn>string</KmsKeyArn>
      <SseAlgorithm>string</SseAlgorithm>
    </EncryptionConfiguration>
    <RecordExpiration>
```

```

    <Days>integer</Days>
    <Expiration>string</Expiration>
  </RecordExpiration>
</JournalTableConfiguration>
<InventoryTableConfiguration>
  <ConfigurationState>string</ConfigurationState>
  <EncryptionConfiguration>
    <KmsKeyArn>string</KmsKeyArn>
    <SseAlgorithm>string</SseAlgorithm>
  </EncryptionConfiguration>
</InventoryTableConfiguration>
<AnnotationTableConfiguration>
  <ConfigurationState>string</ConfigurationState>
  <EncryptionConfiguration>
    <KmsKeyArn>string</KmsKeyArn>
    <SseAlgorithm>string</SseAlgorithm>
  </EncryptionConfiguration>
  <Role>string</Role>
</AnnotationTableConfiguration>
</MetadataConfiguration>

```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The general purpose bucket that you want to create the metadata configuration for.

Required: Yes

Content-MD5

The Content-MD5 header for the metadata configuration.

x-amz-expected-bucket-owner

The expected owner of the general purpose bucket that corresponds to your metadata configuration.

x-amz-sdk-checksum-algorithm

The checksum algorithm to use with your metadata configuration.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 |
XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in XML format.

MetadataConfiguration

Root level tag for the MetadataConfiguration parameters.

Required: Yes

AnnotationTableConfiguration

Optional annotation table configuration to include with the metadata configuration.

Type: [AnnotationTableConfiguration](#) data type

Required: No

InventoryTableConfiguration

The inventory table configuration for a metadata configuration.

Type: [InventoryTableConfiguration](#) data type

Required: No

JournalTableConfiguration

The journal table configuration for a metadata configuration.

Type: [JournalTableConfiguration](#) data type

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Trust policy

This example illustrates one usage of `CreateBucketMetadataConfiguration`.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "metadata.s3.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "123456789012"
        },
        "ArnLike": {
          "aws:SourceArn": "arn:aws:s3:::amzn-s3-demo-bucket"
        }
      }
    }
  ]
}
```

Permissions policy

This example illustrates one usage of `CreateBucketMetadataConfiguration`.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "PermissionForGetAnnotation",
      "Effect": "Allow",
      "Action": [
        "s3:GetObjectAnnotation",
        "s3:GetObjectVersionAnnotation"
      ]
    }
  ]
}
```

```

    ],
    "Resource": ["arn:aws:s3:::amzn-s3-demo-bucket/*"],
    "Condition": {
        "StringEquals": {
            "aws:ResourceAccount": "123456789012"
        }
    }
},
{
    "Sid": "PermissionsForListBucket",
    "Effect": "Allow",
    "Action": [
        "s3:ListBucket",
        "s3:ListBucketVersions"
    ],
    "Resource": ["arn:aws:s3:::amzn-s3-demo-bucket"],
    "Condition": {
        "StringEquals": {
            "aws:ResourceAccount": "123456789012"
        }
    }
},
{
    "Sid": "PermissionsForDecryptAnnotation",
    "Effect": "Allow",
    "Action": ["kms:Decrypt"],
    "Condition": {
        "StringLike": {
            "kms:ViaService": [
                "s3.us-east-1.amazonaws.com"
            ]
        }
    },
    "ArnLike": {
        "kms:EncryptionContext:aws:s3:arn": [
            "arn:aws:s3:::amzn-s3-demo-bucket",
            "arn:aws:s3:::amzn-s3-demo-bucket/*"
        ]
    }
},
{
    "Resource": ["arn:aws:kms:us-east-1:123456789012:key/1234abcd-12ab-34cd-56ef-1234567890ab"]
}
]
}

```


See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

CreateBucketMetadataTableConfiguration

Service: Amazon S3

Important

We recommend that you create your S3 Metadata configurations by using the V2 [CreateBucketMetadataConfiguration](#) API operation. We no longer recommend using the V1 `CreateBucketMetadataTableConfiguration` API operation.

If you created your S3 Metadata configuration before July 15, 2025, we recommend that you delete and re-create your configuration by using [CreateBucketMetadataConfiguration](#) so that you can expire journal table records and create a live inventory table.

Creates a V1 S3 Metadata configuration for a general purpose bucket. For more information, see [Accelerating data discovery with S3 Metadata](#) in the *Amazon S3 User Guide*.

Permissions

To use this operation, you must have the following permissions. For more information, see [Setting up permissions for configuring metadata tables](#) in the *Amazon S3 User Guide*.

If you want to encrypt your metadata tables with server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS), you need additional permissions. For more information, see [Setting up permissions for configuring metadata tables](#) in the *Amazon S3 User Guide*.

If you also want to integrate your table bucket with AWS analytics services so that you can query your metadata table, you need additional permissions. For more information, see [Integrating Amazon S3 Tables with AWS analytics services](#) in the *Amazon S3 User Guide*.

- `s3:CreateBucketMetadataTableConfiguration`
- `s3tables:CreateNamespace`
- `s3tables:GetTable`
- `s3tables:CreateTable`
- `s3tables:PutTablePolicy`

The following operations are related to `CreateBucketMetadataTableConfiguration`:

- [DeleteBucketMetadataTableConfiguration](#)
- [GetBucketMetadataTableConfiguration](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
POST /?metadataTable HTTP/1.1
Host: Bucket.s3.amazonaws.com
Content-MD5: ContentMD5
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<MetadataTableConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <S3TablesDestination>
    <TableBucketArn>string</TableBucketArn>
    <TableName>string</TableName>
  </S3TablesDestination>
</MetadataTableConfiguration>
```

URI Request Parameters

The request uses the following URI parameters.

[Bucket](#)

The general purpose bucket that you want to create the metadata table configuration for.

Required: Yes

[Content-MD5](#)

The Content-MD5 header for the metadata table configuration.

[x-amz-expected-bucket-owner](#)

The expected owner of the general purpose bucket that corresponds to your metadata table configuration.

[x-amz-sdk-checksum-algorithm](#)

The checksum algorithm to use with your metadata table configuration.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 | XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in XML format.

[MetadataTableConfiguration](#)

Root level tag for the MetadataTableConfiguration parameters.

Required: Yes

[S3TablesDestination](#)

The destination information for the metadata table configuration. The destination table bucket must be in the same Region and AWS account as the general purpose bucket. The specified metadata table name must be unique within the `aws_s3_metadata` namespace in the destination table bucket.

Type: [S3TablesDestination](#) data type

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

CreateMultipartUpload

Service: Amazon S3

Important

End of support notice: As of October 1, 2025, Amazon S3 has discontinued support for Email Grantee Access Control Lists (ACLs). If you attempt to use an Email Grantee ACL in a request after October 1, 2025, the request will receive an HTTP 405 (Method Not Allowed) error.

This change affects the following AWS Regions: US East (N. Virginia), US West (N. California), US West (Oregon), Asia Pacific (Singapore), Asia Pacific (Sydney), Asia Pacific (Tokyo), Europe (Ireland), and South America (São Paulo).

This action initiates a multipart upload and returns an upload ID. This upload ID is used to associate all of the parts in the specific multipart upload. You specify this upload ID in each of your subsequent upload part requests (see [UploadPart](#)). You also include this upload ID in the final request to either complete or abort the multipart upload request. For more information about multipart uploads, see [Multipart Upload Overview](#) in the *Amazon S3 User Guide*.

Note

After you initiate a multipart upload and upload one or more parts, to stop being charged for storing the uploaded parts, you must either complete or abort the multipart upload. Amazon S3 frees up the space used to store the parts and stops charging you for storing them only after you either complete or abort a multipart upload.

If you have configured a lifecycle rule to abort incomplete multipart uploads, the created multipart upload must be completed within the number of days specified in the bucket lifecycle configuration. Otherwise, the incomplete multipart upload becomes eligible for an abort action and Amazon S3 aborts the multipart upload. For more information, see [Aborting Incomplete Multipart Uploads Using a Bucket Lifecycle Configuration](#).

Note

- **Directory buckets** - S3 Lifecycle is not supported by directory buckets.

- **Directory buckets** - For directory buckets, you must make requests for this API operation to the Zonal endpoint. These endpoints support virtual-hosted-style requests in the format `https://amzn-s3-demo-bucket.s3express-zone-id.region-code.amazonaws.com/key-name`. Path-style requests are not supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

Request signing

For request signing, multipart upload is just a series of regular requests. You initiate a multipart upload, send one or more requests to upload parts, and then complete the multipart upload process. You sign each request individually. There is nothing special about signing multipart upload requests. For more information about signing, see [Authenticating Requests \(AWS Signature Version 4\)](#) in the *Amazon S3 User Guide*.

Permissions

- **General purpose bucket permissions** - To perform a multipart upload with encryption using an AWS Key Management Service (AWS KMS) KMS key, the requester must have permission to the `kms:Decrypt` and `kms:GenerateDataKey` actions on the key. The requester must also have permissions for the `kms:GenerateDataKey` action for the `CreateMultipartUpload` API. Then, the requester needs permissions for the `kms:Decrypt` action on the `UploadPart` and `UploadPartCopy` APIs. These permissions are required because Amazon S3 must decrypt and read data from the encrypted file parts before it completes the multipart upload. For more information, see [Multipart upload API and permissions](#) and [Protecting data using server-side encryption with AWS KMS](#) in the *Amazon S3 User Guide*.
- **Directory bucket permissions** - To grant access to this API operation on a directory bucket, we recommend that you use the [CreateSession](#) API operation for session-based authorization. Specifically, you grant the `s3express:CreateSession` permission to the directory bucket in a bucket policy or an IAM identity-based policy. Then, you make the `CreateSession` API call on the bucket to obtain a session token. With the session token in your request header, you can make API requests to this operation. After the session token expires, you make another `CreateSession` API call to generate a new session token for use. AWS CLI or SDKs create session and refresh the session token automatically to avoid

service interruptions when a session expires. For more information about authorization, see [CreateSession](#).

Encryption

- **General purpose buckets** - Server-side encryption is for data encryption at rest. Amazon S3 encrypts your data as it writes it to disks in its data centers and decrypts it when you access it. Amazon S3 automatically encrypts all new objects that are uploaded to an S3 bucket. When doing a multipart upload, if you don't specify encryption information in your request, the encryption setting of the uploaded parts is set to the default encryption configuration of the destination bucket. By default, all buckets have a base level of encryption configuration that uses server-side encryption with Amazon S3 managed keys (SSE-S3). If the destination bucket has a default encryption configuration that uses server-side encryption with an AWS Key Management Service (AWS KMS) key (SSE-KMS), or a customer-provided encryption key (SSE-C), Amazon S3 uses the corresponding KMS key, or a customer-provided key to encrypt the uploaded parts. When you perform a `CreateMultipartUpload` operation, if you want to use a different type of encryption setting for the uploaded parts, you can request that Amazon S3 encrypts the object with a different encryption key (such as an Amazon S3 managed key, a KMS key, or a customer-provided key). When the encryption setting in your request is different from the default encryption configuration of the destination bucket, the encryption setting in your request takes precedence. If you choose to provide your own encryption key, the request headers you provide in [UploadPart](#) and [UploadPartCopy](#) requests must match the headers you used in the `CreateMultipartUpload` request.
- Use KMS keys (SSE-KMS) that include the AWS managed key (`aws/s3`) and AWS KMS customer managed keys stored in AWS Key Management Service (AWS KMS) – If you want AWS to manage the keys used to encrypt data, specify the following headers in the request.
 - `x-amz-server-side-encryption`
 - `x-amz-server-side-encryption-aws-kms-key-id`
 - `x-amz-server-side-encryption-context`

Note

- If you specify `x-amz-server-side-encryption:aws:kms`, but don't provide `x-amz-server-side-encryption-aws-kms-key-id`, Amazon S3 uses the AWS managed key (`aws/s3` key) in AWS KMS to protect the data.
- To perform a multipart upload with encryption by using an AWS KMS key, the requester must have permission to the `kms:Decrypt` and

`kms:GenerateDataKey*` actions on the key. These permissions are required because Amazon S3 must decrypt and read data from the encrypted file parts before it completes the multipart upload. For more information, see [Multipart upload API and permissions](#) and [Protecting data using server-side encryption with AWS KMS](#) in the *Amazon S3 User Guide*.

- If your AWS Identity and Access Management (IAM) user or role is in the same AWS account as the KMS key, then you must have these permissions on the key policy. If your IAM user or role is in a different account from the key, then you must have the permissions on both the key policy and your IAM user or role.
- All GET and PUT requests for an object protected by AWS KMS fail if you don't make them by using Secure Sockets Layer (SSL), Transport Layer Security (TLS), or Signature Version 4. For information about configuring any of the officially supported AWS SDKs and AWS CLI, see [Specifying the Signature Version in Request Authentication](#) in the *Amazon S3 User Guide*.

For more information about server-side encryption with AWS KMS keys (SSE-KMS), see [Protecting Data Using Server-Side Encryption with KMS keys](#) in the *Amazon S3 User Guide*.

- Use customer-provided encryption keys (SSE-C) – If you want to manage your own encryption keys, provide all the following headers in the request.
 - `x-amz-server-side-encryption-customer-algorithm`
 - `x-amz-server-side-encryption-customer-key`
 - `x-amz-server-side-encryption-customer-key-MD5`

For more information about server-side encryption with customer-provided encryption keys (SSE-C), see [Protecting data using server-side encryption with customer-provided encryption keys \(SSE-C\)](#) in the *Amazon S3 User Guide*.

- **Directory buckets** - For directory buckets, there are only two supported options for server-side encryption: server-side encryption with Amazon S3 managed keys (SSE-S3) (AES256) and server-side encryption with AWS KMS keys (SSE-KMS) (`aws:kms`). We recommend that the bucket's default encryption uses the desired encryption configuration and you don't override the bucket default encryption in your `CreateSession` requests or PUT object requests. Then, new objects are automatically encrypted with the desired encryption settings. For more information, see [Protecting data with server-side encryption](#) in the *Amazon S3 User*

Guide. For more information about the encryption overriding behaviors in directory buckets, see [Specifying server-side encryption with AWS KMS for new object uploads](#).

In the Zonal endpoint API calls (except [CopyObject](#) and [UploadPartCopy](#)) using the REST API, the encryption request headers must match the encryption settings that are specified in the `CreateSession` request. You can't override the values of the encryption settings (`x-amz-server-side-encryption`, `x-amz-server-side-encryption-aws-kms-key-id`, `x-amz-server-side-encryption-context`, and `x-amz-server-side-encryption-bucket-key-enabled`) that are specified in the `CreateSession` request. You don't need to explicitly specify these encryption settings values in Zonal endpoint API calls, and Amazon S3 will use the encryption settings values from the `CreateSession` request to protect new objects in the directory bucket.

 Note

When you use the CLI or the AWS SDKs, for `CreateSession`, the session token refreshes automatically to avoid service interruptions when a session expires. The CLI or the AWS SDKs use the bucket's default encryption configuration for the `CreateSession` request. It's not supported to override the encryption settings values in the `CreateSession` request. So in the Zonal endpoint API calls (except [CopyObject](#) and [UploadPartCopy](#)), the encryption request headers must match the default encryption configuration of the directory bucket.

 Note

For directory buckets, when you perform a `CreateMultipartUpload` operation and an `UploadPartCopy` operation, the request headers you provide in the `CreateMultipartUpload` request must match the default encryption configuration of the destination bucket.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `Bucket-name.s3express-zone-id.region-code.amazonaws.com`.

The following operations are related to `CreateMultipartUpload`:

- [UploadPart](#)
- [CompleteMultipartUpload](#)
- [AbortMultipartUpload](#)
- [ListParts](#)
- [ListMultipartUploads](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
POST /{Key+}?uploads HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-acl: ACL
Cache-Control: CacheControl
Content-Disposition: ContentDisposition
Content-Encoding: ContentEncoding
Content-Language: ContentLanguage
Content-Type: ContentType
Expires: Expires
x-amz-grant-full-control: GrantFullControl
x-amz-grant-read: GrantRead
x-amz-grant-read-acp: GrantReadACP
x-amz-grant-write-acp: GrantWriteACP
x-amz-server-side-encryption: ServerSideEncryption
x-amz-storage-class: StorageClass
x-amz-website-redirect-location: WebsiteRedirectLocation
x-amz-server-side-encryption-customer-algorithm: SSECustomerAlgorithm
x-amz-server-side-encryption-customer-key: SSECustomerKey
x-amz-server-side-encryption-customer-key-MD5: SSECustomerKeyMD5
x-amz-server-side-encryption-aws-kms-key-id: SSEKMSKeyId
x-amz-server-side-encryption-context: SSEKMSEncryptionContext
x-amz-server-side-encryption-bucket-key-enabled: BucketKeyEnabled
x-amz-request-payer: RequestPayer
x-amz-tagging: Tagging
x-amz-object-lock-mode: ObjectLockMode
```

```
x-amz-object-lock-retain-until-date: ObjectLockRetainUntilDate  
x-amz-object-lock-legal-hold: ObjectLockLegalHoldStatus  
x-amz-object-lock-event-hold: ObjectLockEventHold  
x-amz-object-lock-event-hold-duration-days: ObjectLockEventHoldDurationDays  
x-amz-object-lock-event-hold-duration-years: ObjectLockEventHoldDurationYears  
x-amz-expected-bucket-owner: ExpectedBucketOwner  
x-amz-checksum-algorithm: ChecksumAlgorithm  
x-amz-checksum-type: ChecksumType
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket where the multipart upload is initiated and where the object is uploaded.

Directory buckets - When you use this operation with a directory bucket, you must use virtual-hosted-style requests in the format *Bucket-name.s3express-zone-id.region-code.amazonaws.com*. Path-style requests are not supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must follow the format *bucket-base-name--zone-id--x-s3* (for example, *amzn-s3-demo-bucket--usw2-az1--x-s3*). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form *AccessPointName-AccountId.s3-accesspoint.Region.amazonaws.com*. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

Note

Object Lambda access points are not supported by directory buckets.

S3 on Outposts - When you use this action with S3 on Outposts, you must direct requests to the S3 on Outposts hostname. The S3 on Outposts hostname takes the form *AccessPointName-AccountId.outpostID.s3-outposts.Region.amazonaws.com*. When you use this action with S3 on Outposts, the destination bucket must be the Outposts access point ARN or the access point alias. For more information about S3 on Outposts, see [What is S3 on Outposts?](#) in the *Amazon S3 User Guide*.

Required: Yes

Cache-Control

Specifies caching behavior along the request/reply chain.

Content-Disposition

Specifies presentational information for the object.

Content-Encoding

Specifies what content encodings have been applied to the object and thus what decoding mechanisms must be applied to obtain the media-type referenced by the Content-Type header field.

Note

For directory buckets, only the `aws-chunked` value is supported in this header field.

Content-Language

The language that the content is in.

Content-Type

A standard MIME type describing the format of the object data.

Expires

The date and time at which the object is no longer cacheable.

Key

Object key for which the multipart upload is to be initiated.

Length Constraints: Minimum length of 1.

Required: Yes

x-amz-acl

The canned ACL to apply to the object. Amazon S3 supports a set of predefined ACLs, known as *canned ACLs*. Each canned ACL has a predefined set of grantees and permissions. For more information, see [Canned ACL](#) in the *Amazon S3 User Guide*.

By default, all objects are private. Only the owner has full access control. When uploading an object, you can grant access permissions to individual AWS accounts or to predefined groups defined by Amazon S3. These permissions are then added to the access control list (ACL) on the new object. For more information, see [Using ACLs](#). One way to grant the permissions using the request headers is to specify a canned ACL with the `x-amz-acl` request header.

Note

- This functionality is not supported for directory buckets.
- This functionality is not supported for Amazon S3 on Outposts.

Valid Values: `private` | `public-read` | `public-read-write` | `authenticated-read` | `aws-exec-read` | `bucket-owner-read` | `bucket-owner-full-control`

x-amz-checksum-algorithm

Indicates the algorithm that you want Amazon S3 to use to create the checksum for the object. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Valid Values: `CRC32` | `CRC32C` | `SHA1` | `SHA256` | `CRC64NVME` | `SHA512` | `MD5` | `XXHASH64` | `XXHASH3` | `XXHASH128`

x-amz-checksum-type

Indicates the checksum type that you want Amazon S3 to use to calculate the object's checksum value. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

Valid Values: `COMPOSITE` | `FULL_OBJECT`

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code `403 Forbidden` (access denied).

x-amz-grant-full-control

Specify access permissions explicitly to give the grantee READ, READ_ACP, and WRITE_ACP permissions on the object.

By default, all objects are private. Only the owner has full access control. When uploading an object, you can use this header to explicitly grant access permissions to specific AWS accounts or groups. This header maps to specific permissions that Amazon S3 supports in an ACL. For more information, see [Access Control List \(ACL\) Overview](#) in the *Amazon S3 User Guide*.

You specify each grantee as a type=value pair, where the type is one of the following:

- `id` – if the value specified is the canonical user ID of an AWS account
- `uri` – if you are granting permissions to a predefined group
- `emailAddress` – if the value specified is the email address of an AWS account

Note

Using email addresses to specify a grantee is only supported in the following AWS Regions:

- US East (N. Virginia)
- US West (N. California)
- US West (Oregon)
- Asia Pacific (Singapore)
- Asia Pacific (Sydney)
- Asia Pacific (Tokyo)
- Europe (Ireland)
- South America (São Paulo)

For a list of all the Amazon S3 supported Regions and endpoints, see [Regions and Endpoints](#) in the AWS General Reference.

For example, the following `x-amz-grant-read` header grants the AWS accounts identified by account IDs permissions to read object data and its metadata:

```
x-amz-grant-read: id="11112222333", id="444455556666"
```

Note

- This functionality is not supported for directory buckets.
- This functionality is not supported for Amazon S3 on Outposts.

x-amz-grant-read

Specify access permissions explicitly to allow grantee to read the object data and its metadata.

By default, all objects are private. Only the owner has full access control. When uploading an object, you can use this header to explicitly grant access permissions to specific AWS accounts or groups. This header maps to specific permissions that Amazon S3 supports in an ACL. For more information, see [Access Control List \(ACL\) Overview](#) in the *Amazon S3 User Guide*.

You specify each grantee as a type=value pair, where the type is one of the following:

- `id` – if the value specified is the canonical user ID of an AWS account
- `uri` – if you are granting permissions to a predefined group
- `emailAddress` – if the value specified is the email address of an AWS account

Note

Using email addresses to specify a grantee is only supported in the following AWS Regions:

- US East (N. Virginia)
- US West (N. California)
- US West (Oregon)
- Asia Pacific (Singapore)
- Asia Pacific (Sydney)
- Asia Pacific (Tokyo)
- Europe (Ireland)
- South America (São Paulo)

For a list of all the Amazon S3 supported Regions and endpoints, see [Regions and Endpoints](#) in the AWS General Reference.

For example, the following `x-amz-grant-read` header grants the AWS accounts identified by account IDs permissions to read object data and its metadata:

```
x-amz-grant-read: id="11112222333", id="444455556666"
```

 Note

- This functionality is not supported for directory buckets.
- This functionality is not supported for Amazon S3 on Outposts.

`x-amz-grant-read-acp`

Specify access permissions explicitly to allow grantee to read the object ACL.

By default, all objects are private. Only the owner has full access control. When uploading an object, you can use this header to explicitly grant access permissions to specific AWS accounts or groups. This header maps to specific permissions that Amazon S3 supports in an ACL. For more information, see [Access Control List \(ACL\) Overview](#) in the *Amazon S3 User Guide*.

You specify each grantee as a `type=value` pair, where the type is one of the following:

- `id` – if the value specified is the canonical user ID of an AWS account
- `uri` – if you are granting permissions to a predefined group
- `emailAddress` – if the value specified is the email address of an AWS account

 Note

Using email addresses to specify a grantee is only supported in the following AWS Regions:

- US East (N. Virginia)
- US West (N. California)
- US West (Oregon)
- Asia Pacific (Singapore)
- Asia Pacific (Sydney)
- Asia Pacific (Tokyo)
- Europe (Ireland)
- South America (São Paulo)

For a list of all the Amazon S3 supported Regions and endpoints, see [Regions and Endpoints](#) in the AWS General Reference.

For example, the following `x-amz-grant-read` header grants the AWS accounts identified by account IDs permissions to read object data and its metadata:

```
x-amz-grant-read: id="11112222333", id="444455556666"
```

Note

- This functionality is not supported for directory buckets.
- This functionality is not supported for Amazon S3 on Outposts.

[x-amz-grant-write-acp](#)

Specify access permissions explicitly to allow grantee to write the ACL for the applicable object.

By default, all objects are private. Only the owner has full access control. When uploading an object, you can use this header to explicitly grant access permissions to specific AWS accounts or groups. This header maps to specific permissions that Amazon S3 supports in an ACL. For more information, see [Access Control List \(ACL\) Overview](#) in the *Amazon S3 User Guide*.

You specify each grantee as a `type=value` pair, where the type is one of the following:

- `id` – if the value specified is the canonical user ID of an AWS account
- `uri` – if you are granting permissions to a predefined group
- `emailAddress` – if the value specified is the email address of an AWS account

Note

Using email addresses to specify a grantee is only supported in the following AWS Regions:

- US East (N. Virginia)
- US West (N. California)
- US West (Oregon)
- Asia Pacific (Singapore)

- Asia Pacific (Sydney)
- Asia Pacific (Tokyo)
- Europe (Ireland)
- South America (São Paulo)

For a list of all the Amazon S3 supported Regions and endpoints, see [Regions and Endpoints](#) in the AWS General Reference.

For example, the following `x-amz-grant-read` header grants the AWS accounts identified by account IDs permissions to read object data and its metadata:

```
x-amz-grant-read: id="11112222333", id="444455556666"
```

 Note

- This functionality is not supported for directory buckets.
- This functionality is not supported for Amazon S3 on Outposts.

[x-amz-object-lock-event-hold](#)

Specifies the event hold status to apply to the uploaded object. Set to ON to enable or OFF to disable.

 Note

This functionality is not supported for directory buckets.

Valid Values: ON | OFF

[x-amz-object-lock-event-hold-duration-days](#)

Specifies the event hold duration in days to apply to the uploaded object. You cannot specify a duration in both days and years.

 Note

This functionality is not supported for directory buckets.

x-amz-object-lock-event-hold-duration-years

Specifies the event hold duration in years to apply to the uploaded object. You cannot specify a duration in both days and years.

Note

This functionality is not supported for directory buckets.

x-amz-object-lock-legal-hold

Specifies whether you want to apply a legal hold to the uploaded object.

Note

This functionality is not supported for directory buckets.

Valid Values: ON | OFF

x-amz-object-lock-mode

Specifies the Object Lock mode that you want to apply to the uploaded object.

Note

This functionality is not supported for directory buckets.

Valid Values: GOVERNANCE | COMPLIANCE

x-amz-object-lock-retain-until-date

Specifies the date and time when you want the Object Lock to expire.

Note

This functionality is not supported for directory buckets.

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

x-amz-server-side-encryption

The server-side encryption algorithm used when you store this object in Amazon S3 or Amazon FSx.

- **Directory buckets** - For directory buckets, there are only two supported options for server-side encryption: server-side encryption with Amazon S3 managed keys (SSE-S3) (AES256) and server-side encryption with AWS KMS keys (SSE-KMS) (aws:kms). We recommend that the bucket's default encryption uses the desired encryption configuration and you don't override the bucket default encryption in your CreateSession requests or PUT object requests. Then, new objects are automatically encrypted with the desired encryption settings. For more information, see [Protecting data with server-side encryption](#) in the *Amazon S3 User Guide*. For more information about the encryption overriding behaviors in directory buckets, see [Specifying server-side encryption with AWS KMS for new object uploads](#).

In the Zonal endpoint API calls (except [CopyObject](#) and [UploadPartCopy](#)) using the REST API, the encryption request headers must match the encryption settings that are specified in the CreateSession request. You can't override the values of the encryption settings (x-amz-server-side-encryption, x-amz-server-side-encryption-aws-kms-key-id, x-amz-server-side-encryption-context, and x-amz-server-side-encryption-bucket-key-enabled) that are specified in the CreateSession request. You don't need to explicitly specify these encryption settings values in Zonal endpoint API calls, and Amazon S3 will use the encryption settings values from the CreateSession request to protect new objects in the directory bucket.

Note

When you use the CLI or the AWS SDKs, for `CreateSession`, the session token refreshes automatically to avoid service interruptions when a session expires. The CLI or the AWS SDKs use the bucket's default encryption configuration for the `CreateSession` request. It's not supported to override the encryption settings values in the `CreateSession` request. So in the Zonal endpoint API calls (except [CopyObject](#) and [UploadPartCopy](#)), the encryption request headers must match the default encryption configuration of the directory bucket.

- **S3 access points for Amazon FSx** - When accessing data stored in Amazon FSx file systems using S3 access points, the only valid server side encryption option is `aws:fsx`. All Amazon FSx file systems have encryption configured by default and are encrypted at rest. Data is automatically encrypted before being written to the file system, and automatically decrypted as it is read. These processes are handled transparently by Amazon FSx.

Valid Values: `AES256` | `aws:fsx` | `aws:backup` | `aws:kms` | `aws:kms:dsse`

`x-amz-server-side-encryption-aws-kms-key-id`

Specifies the AWS KMS key ID (Key ID, Key ARN, or Key Alias) to use for object encryption. If the KMS key doesn't exist in the same account that's issuing the command, you must use the full Key ARN not the Key ID.

General purpose buckets - If you specify `x-amz-server-side-encryption` with `aws:kms` or `aws:kms:dsse`, this header specifies the ID (Key ID, Key ARN, or Key Alias) of the AWS KMS key to use. If you specify `x-amz-server-side-encryption:aws:kms` or `x-amz-server-side-encryption:aws:kms:dsse`, but do not provide `x-amz-server-side-encryption-aws-kms-key-id`, Amazon S3 uses the AWS managed key (`aws/s3`) to protect the data.

Directory buckets - To encrypt data using SSE-KMS, it's recommended to specify the `x-amz-server-side-encryption` header to `aws:kms`. Then, the `x-amz-server-side-encryption-aws-kms-key-id` header implicitly uses the bucket's default KMS customer managed key ID. If you want to explicitly set the `x-amz-server-side-encryption-aws-kms-key-id` header, it must match the bucket's default customer managed key (using key ID or ARN, not alias). Your SSE-KMS configuration can only support 1 [customer managed key](#) per directory bucket's lifetime. The [AWS managed key](#) (`aws/s3`) isn't supported. Incorrect key specification results in an HTTP 400 Bad Request error.

x-amz-server-side-encryption-bucket-key-enabled

Specifies whether Amazon S3 should use an S3 Bucket Key for object encryption with server-side encryption using AWS Key Management Service (AWS KMS) keys (SSE-KMS).

General purpose buckets - Setting this header to `true` causes Amazon S3 to use an S3 Bucket Key for object encryption with SSE-KMS. Also, specifying this header with a PUT action doesn't affect bucket-level settings for S3 Bucket Key.

Directory buckets - S3 Bucket Keys are always enabled for GET and PUT operations in a directory bucket and can't be disabled. S3 Bucket Keys aren't supported, when you copy SSE-KMS encrypted objects from general purpose buckets to directory buckets, from directory buckets to general purpose buckets, or between directory buckets, through [CopyObject](#), [UploadPartCopy](#), [the Copy operation in Batch Operations](#), or [the import jobs](#). In this case, Amazon S3 makes a call to AWS KMS every time a copy request is made for a KMS-encrypted object.

x-amz-server-side-encryption-context

Specifies the AWS KMS Encryption Context to use for object encryption. The value of this header is a Base64 encoded string of a UTF-8 encoded JSON, which contains the encryption context as key-value pairs.

Directory buckets - You can optionally provide an explicit encryption context value. The value must match the default encryption context - the bucket Amazon Resource Name (ARN). An additional encryption context value is not supported.

x-amz-server-side-encryption-customer-algorithm

Specifies the algorithm to use when encrypting the object (for example, AES256).

Note

This functionality is not supported for directory buckets.

x-amz-server-side-encryption-customer-key

Specifies the customer-provided encryption key for Amazon S3 to use in encrypting data. This value is used to store the object and then it is discarded; Amazon S3 does not store the

encryption key. The key must be appropriate for use with the algorithm specified in the `x-amz-server-side-encryption-customer-algorithm` header.

Note

This functionality is not supported for directory buckets.

x-amz-server-side-encryption-customer-key-MD5

Specifies the 128-bit MD5 digest of the customer-provided encryption key according to RFC 1321. Amazon S3 uses this header for a message integrity check to ensure that the encryption key was transmitted without error.

Note

This functionality is not supported for directory buckets.

x-amz-storage-class

By default, Amazon S3 uses the STANDARD Storage Class to store newly created objects. The STANDARD storage class provides high durability and high availability. Depending on performance needs, you can specify a different Storage Class. For more information, see [Storage Classes](#) in the *Amazon S3 User Guide*.

Note

- Directory buckets only support EXPRESS_ONEZONE (the S3 Express One Zone storage class) in Availability Zones and ONEZONE_IA (the S3 One Zone-Infrequent Access storage class) in Dedicated Local Zones.
- Amazon S3 on Outposts only uses the OUTPOSTS Storage Class.

Valid Values: STANDARD | REDUCED_REDUNDANCY | STANDARD_IA | ONEZONE_IA | INTELLIGENT_TIERING | GLACIER | DEEP_ARCHIVE | OUTPOSTS | GLACIER_IR | SNOW | EXPRESS_ONEZONE | FSX_OPENZFS | FSX_ONTAP | AWS_BACKUP_WARM | AWS_BACKUP_LOW_COST_WARM

x-amz-tagging

The tag-set for the object. The tag-set must be encoded as URL Query parameters.

Note

This functionality is not supported for directory buckets.

x-amz-website-redirect-location

If the bucket is configured as a website, redirects requests for this object to another object in the same bucket or to an external URL. Amazon S3 stores the value of this header in the object metadata.

Note

This functionality is not supported for directory buckets.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
x-amz-abort-date: AbortDate
x-amz-abort-rule-id: AbortRuleId
x-amz-server-side-encryption: ServerSideEncryption
x-amz-server-side-encryption-customer-algorithm: SSECustomerAlgorithm
x-amz-server-side-encryption-customer-key-MD5: SSECustomerKeyMD5
x-amz-server-side-encryption-aws-kms-key-id: SSEKMSKeyId
x-amz-server-side-encryption-context: SSEKMSEncryptionContext
x-amz-server-side-encryption-bucket-key-enabled: BucketKeyEnabled
x-amz-request-charged: RequestCharged
x-amz-checksum-algorithm: ChecksumAlgorithm
x-amz-checksum-type: ChecksumType
<?xml version="1.0" encoding="UTF-8"?>
<InitiateMultipartUploadResult>
  <Bucket>string</Bucket>
  <Key>string</Key>
```

```
<UploadId>string</UploadId>  
</InitiateMultipartUploadResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

x-amz-abort-date

If the bucket has a lifecycle rule configured with an action to abort incomplete multipart uploads and the prefix in the lifecycle rule matches the object name in the request, the response includes this header. The header indicates when the initiated multipart upload becomes eligible for an abort operation. For more information, see [Aborting Incomplete Multipart Uploads Using a Bucket Lifecycle Configuration](#) in the *Amazon S3 User Guide*.

The response also includes the `x-amz-abort-rule-id` header that provides the ID of the lifecycle configuration rule that defines the abort action.

Note

This functionality is not supported for directory buckets.

x-amz-abort-rule-id

This header is returned along with the `x-amz-abort-date` header. It identifies the applicable lifecycle configuration rule that defines the action to abort incomplete multipart uploads.

Note

This functionality is not supported for directory buckets.

x-amz-checksum-algorithm

The algorithm that was used to create a checksum of the object.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 |
XXHASH64 | XXHASH3 | XXHASH128

x-amz-checksum-type

Indicates the checksum type that you want Amazon S3 to use to calculate the object's checksum value. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

Valid Values: COMPOSITE | FULL_OBJECT

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

x-amz-server-side-encryption

The server-side encryption algorithm used when you store this object in Amazon S3 or Amazon FSx.

Note

When accessing data stored in Amazon FSx file systems using S3 access points, the only valid server side encryption option is `aws:fsx`.

Valid Values: AES256 | aws:fsx | aws:backup | aws:kms | aws:kms:dsse

x-amz-server-side-encryption-aws-kms-key-id

If present, indicates the ID of the KMS key that was used for object encryption.

x-amz-server-side-encryption-bucket-key-enabled

Indicates whether the multipart upload uses an S3 Bucket Key for server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS).

x-amz-server-side-encryption-context

If present, indicates the AWS KMS Encryption Context to use for object encryption. The value of this header is a Base64 encoded string of a UTF-8 encoded JSON, which contains the encryption context as key-value pairs.

x-amz-server-side-encryption-customer-algorithm

If server-side encryption with a customer-provided encryption key was requested, the response will include this header to confirm the encryption algorithm that's used.

Note

This functionality is not supported for directory buckets.

x-amz-server-side-encryption-customer-key-MD5

If server-side encryption with a customer-provided encryption key was requested, the response will include this header to provide the round-trip message integrity verification of the customer-provided encryption key.

Note

This functionality is not supported for directory buckets.

The following data is returned in XML format by the service.

InitiateMultipartUploadResult

Root level tag for the InitiateMultipartUploadResult parameters.

Required: Yes

Bucket

The name of the bucket to which the multipart upload was initiated. Does not return the access point ARN or access point alias if used.

Note

Access points are not supported by directory buckets.

Type: String

Key

Object key for which the multipart upload was initiated.

Type: String

Length Constraints: Minimum length of 1.

UploadId

ID for the initiated multipart upload.

Type: String

Examples**Sample Request for general purpose buckets**

This action initiates a multipart upload for the `amzn-s3-demo-bucket` object.

```
POST /example-object?uploads HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Mon, 1 Nov 2010 20:34:56 GMT
Authorization: authorization string
```

Sample Response for general purpose buckets

This example illustrates one usage of `CreateMultipartUpload`.

```
HTTP/1.1 200 OK
x-amz-id-2: Uuag1LuByRx9e6j50nimru9p04ZVKnJ2Qz7/C1NPcfTWAtRPfTa0Fg==
```

```

x-amz-request-id: 656c76696e6727732072657175657374
Date: Mon, 1 Nov 2010 20:34:56 GMT
Transfer-Encoding: chunked
Connection: keep-alive
Server: AmazonS3

<?xml version="1.0" encoding="UTF-8"?>
<InitiateMultipartUploadResult xmlns="http://s3.amazonaws.com/
doc/2006-03-01/">
  <Bucket>amzn-s3-demo-bucket</Bucket>
  <Key>example-object</Key>
  <UploadId>VXBsb2FkIElEIGZvciA2aWpmbmcncyBteS1tb3ZpZS5tMnRzIHVwbG9hZA</
UploadId>
</InitiateMultipartUploadResult>

```

Example for general purpose buckets: Initiate a multipart upload using server-side encryption with customer-provided encryption keys

This example, which initiates a multipart upload request, specifies server-side encryption with customer-provided encryption keys by adding relevant headers.

Note

If you have server-side encryption with customer-provided keys (SSE-C) blocked for your general purpose bucket, you will get an HTTP 403 Access Denied error when you specify the SSE-C request headers while writing new data to your bucket. For more information, see [Blocking or unblocking SSE-C for a general purpose bucket](#).

```

POST /example-object?uploads HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Authorization: authorization string
Date: Wed, 28 May 2014 19:34:57 +0000
x-amz-server-side-encryption-customer-key:
g0lCfA3Dv40jZz5SQJ1ZukLRFqtI5WorC/8SEEXAMPLE
x-amz-server-side-encryption-customer-key-MD5: ZjQrne1X/iTcskby2example
x-amz-server-side-encryption-customer-algorithm: AES256

```

Sample Response for general purpose buckets

In the response, Amazon S3 returns an `UploadId`. In addition, Amazon S3 returns the encryption algorithm and the MD5 digest of the encryption key that you provided in the request.

```
HTTP/1.1 200 OK
x-amz-id-2:
36HRCaIGp57F1FvWvVRvd3hNn9WoBGfEaCVHTCt8QWf00qxdHazQUgfoXAbhFWD
x-amz-request-id: 50FA1D691B62CA43
Date: Wed, 28 May 2014 19:34:58 GMT
x-amz-server-side-encryption-customer-algorithm: AES256
x-amz-server-side-encryption-customer-key-MD5: ZjQrne1X/iTcskbY2m3tFg==
Transfer-Encoding: chunked

<?xml version="1.0" encoding="UTF-8"?>
<InitiateMultipartUploadResult
  xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Bucket>amzn-s3-demo-bucket</Bucket>
  <Key>example-object</Key>

  <UploadId>EXAMPLEJZ6e0YupT2h66iePQCc9IEbYbDUy4RTpMeoSMLPRp8Z5o1u8feSRonpvnWsKKG35tI2LB9VDPiCgT
</UploadId>
</InitiateMultipartUploadResult>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)

- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

CreateSession

Service: Amazon S3

Creates a session that establishes temporary security credentials to support fast authentication and authorization for the Zonal endpoint API operations on directory buckets. For more information about Zonal endpoint API operations that include the Availability Zone in the request endpoint, see [S3 Express One Zone APIs](#) in the *Amazon S3 User Guide*.

To make Zonal endpoint API requests on a directory bucket, use the `CreateSession` API operation. Specifically, you grant `s3express:CreateSession` permission to a bucket in a bucket policy or an IAM identity-based policy. Then, you use IAM credentials to make the `CreateSession` API request on the bucket, which returns temporary security credentials that include the access key ID, secret access key, session token, and expiration. These credentials have associated permissions to access the Zonal endpoint API operations. After the session is created, you don't need to use other policies to grant permissions to each Zonal endpoint API individually. Instead, in your Zonal endpoint API requests, you sign your requests by applying the temporary security credentials of the session to the request headers and following the SigV4 protocol for authentication. You also apply the session token to the `x-amz-s3session-token` request header for authorization. Temporary security credentials are scoped to the bucket and expire after 5 minutes. After the expiration time, any calls that you make with those credentials will fail. You must use IAM credentials again to make a `CreateSession` API request that generates a new set of temporary credentials for use. Temporary credentials cannot be extended or refreshed beyond the original specified interval.

If you use AWS SDKs, SDKs handle the session token refreshes automatically to avoid service interruptions when a session expires. We recommend that you use the AWS SDKs to initiate and manage requests to the `CreateSession` API. For more information, see [Performance guidelines and design patterns](#) in the *Amazon S3 User Guide*.

Note

- You must make requests for this API operation to the Zonal endpoint. These endpoints support virtual-hosted-style requests in the format `https://bucket-name.s3express-zone-id.region-code.amazonaws.com`. Path-style requests are not supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

- **CopyObject API operation** - Unlike other Zonal endpoint API operations, the CopyObject API operation doesn't use the temporary security credentials returned from the CreateSession API operation for authentication and authorization. For information about authentication and authorization of the CopyObject API operation on directory buckets, see [CopyObject](#).
- **HeadBucket API operation** - Unlike other Zonal endpoint API operations, the HeadBucket API operation doesn't use the temporary security credentials returned from the CreateSession API operation for authentication and authorization. For information about authentication and authorization of the HeadBucket API operation on directory buckets, see [HeadBucket](#).

Permissions

To obtain temporary security credentials, you must create a bucket policy or an IAM identity-based policy that grants `s3express:CreateSession` permission to the bucket. In a policy, you can have the `s3express:SessionMode` condition key to control who can create a `ReadWrite` or `ReadOnly` session. For more information about `ReadWrite` or `ReadOnly` sessions, see [x-amz-create-session-mode](#). For example policies, see [Example bucket policies for S3 Express One Zone](#) and [AWS Identity and Access Management \(IAM\) identity-based policies for S3 Express One Zone](#) in the *Amazon S3 User Guide*.

To grant cross-account access to Zonal endpoint API operations, the bucket policy should also grant both accounts the `s3express:CreateSession` permission.

If you want to encrypt objects with SSE-KMS, you must also have the `kms:GenerateDataKey` and the `kms:Decrypt` permissions in IAM identity-based policies and AWS KMS key policies for the target AWS KMS key.

Encryption

For directory buckets, there are only two supported options for server-side encryption: server-side encryption with Amazon S3 managed keys (SSE-S3) (AES256) and server-side encryption with AWS KMS keys (SSE-KMS) (`aws:kms`). We recommend that the bucket's default encryption uses the desired encryption configuration and you don't override the bucket default encryption in your `CreateSession` requests or `PUT` object requests. Then, new objects are automatically encrypted with the desired encryption settings. For more information, see [Protecting data with server-side encryption](#) in the *Amazon S3 User Guide*. For more information about the encryption

overriding behaviors in directory buckets, see [Specifying server-side encryption with AWS KMS for new object uploads](#).

For [Zonal endpoint \(object-level\) API operations](#) except [CopyObject](#) and [UploadPartCopy](#), you authenticate and authorize requests through [CreateSession](#) for low latency. To encrypt new objects in a directory bucket with SSE-KMS, you must specify SSE-KMS as the directory bucket's default encryption configuration with a KMS key (specifically, a [customer managed key](#)). Then, when a session is created for Zonal endpoint API operations, new objects are automatically encrypted and decrypted with SSE-KMS and S3 Bucket Keys during the session.

 Note

Only 1 [customer managed key](#) is supported per directory bucket for the lifetime of the bucket. The [AWS managed key](#) (aws/s3) isn't supported. After you specify SSE-KMS as your bucket's default encryption configuration with a customer managed key, you can't change the customer managed key for the bucket's SSE-KMS configuration.

In the Zonal endpoint API calls (except [CopyObject](#) and [UploadPartCopy](#)) using the REST API, you can't override the values of the encryption settings (x-amz-server-side-encryption, x-amz-server-side-encryption-aws-kms-key-id, x-amz-server-side-encryption-context, and x-amz-server-side-encryption-bucket-key-enabled) from the `CreateSession` request. You don't need to explicitly specify these encryption settings values in Zonal endpoint API calls, and Amazon S3 will use the encryption settings values from the `CreateSession` request to protect new objects in the directory bucket.

 Note

When you use the CLI or the AWS SDKs, for `CreateSession`, the session token refreshes automatically to avoid service interruptions when a session expires. The CLI or the AWS SDKs use the bucket's default encryption configuration for the `CreateSession` request. It's not supported to override the encryption settings values in the `CreateSession` request. Also, in the Zonal endpoint API calls (except [CopyObject](#) and [UploadPartCopy](#)), it's not supported to override the values of the encryption settings from the `CreateSession` request.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `Bucket-name.s3express-zone-id.region-code.amazonaws.com`.

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?session HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-create-session-mode: SessionMode
x-amz-server-side-encryption: ServerSideEncryption
x-amz-server-side-encryption-aws-kms-key-id: SSEKMSKeyId
x-amz-server-side-encryption-context: SSEKMSEncryptionContext
x-amz-server-side-encryption-bucket-key-enabled: BucketKeyEnabled
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket that you create a session for.

Required: Yes

x-amz-create-session-mode

Specifies the mode of the session that will be created, either `ReadWrite` or `ReadOnly`. If no session mode is specified, the default behavior attempts to create a session with the maximum allowable privilege. It will first attempt to create a `ReadWrite` session, and if that is not allowed by permissions, it will attempt to create a `ReadOnly` session. If neither session type is allowed, the request will return an Access Denied error. A `ReadWrite` session is capable of executing all the Zonal endpoint API operations on a directory bucket. A `ReadOnly` session is

constrained to execute the following Zonal endpoint API operations: `GetObject`, `HeadObject`, `ListObjectsV2`, `GetObjectAttributes`, `ListParts`, and `ListMultipartUploads`.

Valid Values: `ReadOnly` | `ReadWrite`

[x-amz-server-side-encryption](#)

The server-side encryption algorithm to use when you store objects in the directory bucket.

For directory buckets, there are only two supported options for server-side encryption: server-side encryption with Amazon S3 managed keys (SSE-S3) (AES256) and server-side encryption with AWS KMS keys (SSE-KMS) (`aws:kms`). By default, Amazon S3 encrypts data with SSE-S3. For more information, see [Protecting data with server-side encryption](#) in the *Amazon S3 User Guide*.

S3 access points for Amazon FSx - When accessing data stored in Amazon FSx file systems using S3 access points, the only valid server side encryption option is `aws:fsx`. All Amazon FSx file systems have encryption configured by default and are encrypted at rest. Data is automatically encrypted before being written to the file system, and automatically decrypted as it is read. These processes are handled transparently by Amazon FSx.

Valid Values: `AES256` | `aws:fsx` | `aws:backup` | `aws:kms` | `aws:kms:dsse`

[x-amz-server-side-encryption-aws-kms-key-id](#)

If you specify `x-amz-server-side-encryption` with `aws:kms`, you must specify the `x-amz-server-side-encryption-aws-kms-key-id` header with the ID (Key ID or Key ARN) of the AWS KMS symmetric encryption customer managed key to use. Otherwise, you get an HTTP 400 Bad Request error. Only use the key ID or key ARN. The key alias format of the KMS key isn't supported. Also, if the KMS key doesn't exist in the same account that's issuing the command, you must use the full Key ARN not the Key ID.

Your SSE-KMS configuration can only support 1 [customer managed key](#) per directory bucket's lifetime. The [AWS managed key](#) (`aws/s3`) isn't supported.

[x-amz-server-side-encryption-bucket-key-enabled](#)

Specifies whether Amazon S3 should use an S3 Bucket Key for object encryption with server-side encryption using AWS KMS keys (SSE-KMS).

S3 Bucket Keys are always enabled for GET and PUT operations in a directory bucket and can't be disabled. S3 Bucket Keys aren't supported, when you copy SSE-KMS encrypted objects

from general purpose buckets to directory buckets, from directory buckets to general purpose buckets, or between directory buckets, through [CopyObject](#), [UploadPartCopy](#), [the Copy operation in Batch Operations](#), or [the import jobs](#). In this case, Amazon S3 makes a call to AWS KMS every time a copy request is made for a KMS-encrypted object.

[x-amz-server-side-encryption-context](#)

Specifies the AWS KMS Encryption Context as an additional encryption context to use for object encryption. The value of this header is a Base64 encoded string of a UTF-8 encoded JSON, which contains the encryption context as key-value pairs. This value is stored as object metadata and automatically gets passed on to AWS KMS for future GetObject operations on this object.

General purpose buckets - This value must be explicitly added during CopyObject operations if you want an additional encryption context for your object. For more information, see [Encryption context](#) in the *Amazon S3 User Guide*.

Directory buckets - You can optionally provide an explicit encryption context value. The value must match the default encryption context - the bucket Amazon Resource Name (ARN). An additional encryption context value is not supported.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
x-amz-server-side-encryption: ServerSideEncryption
x-amz-server-side-encryption-aws-kms-key-id: SSEKMSKeyId
x-amz-server-side-encryption-context: SSEKMSEncryptionContext
x-amz-server-side-encryption-bucket-key-enabled: BucketKeyEnabled
<?xml version="1.0" encoding="UTF-8"?>
<CreateSessionResult>
  <Credentials>
    <AccessKeyId>string</AccessKeyId>
    <Expiration>timestamp</Expiration>
    <SecretAccessKey>string</SecretAccessKey>
    <SessionToken>string</SessionToken>
  </Credentials>
</CreateSessionResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

[x-amz-server-side-encryption](#)

The server-side encryption algorithm used when you store objects in the directory bucket.

Note

When accessing data stored in Amazon FSx file systems using S3 access points, the only valid server side encryption option is `aws:fsx`.

Valid Values: AES256 | `aws:fsx` | `aws:backup` | `aws:kms` | `aws:kms:dsse`

[x-amz-server-side-encryption-aws-kms-key-id](#)

If you specify `x-amz-server-side-encryption` with `aws:kms`, this header indicates the ID of the AWS KMS symmetric encryption customer managed key that was used for object encryption.

[x-amz-server-side-encryption-bucket-key-enabled](#)

Indicates whether to use an S3 Bucket Key for server-side encryption with AWS KMS keys (SSE-KMS).

[x-amz-server-side-encryption-context](#)

If present, indicates the AWS KMS Encryption Context to use for object encryption. The value of this header is a Base64 encoded string of a UTF-8 encoded JSON, which contains the encryption context as key-value pairs. This value is stored as object metadata and automatically gets passed on to AWS KMS for future `GetObject` operations on this object.

The following data is returned in XML format by the service.

[CreateSessionResult](#)

Root level tag for the `CreateSessionResult` parameters.

Required: Yes

Credentials

The established temporary security credentials for the created session.

Type: [SessionCredentials](#) data type

Errors

NoSuchBucket

The specified bucket does not exist.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteBucket

Service: Amazon S3

Deletes the S3 bucket. All objects (including all object versions and delete markers) in the bucket must be deleted before the bucket itself can be deleted.

Note

- **Directory buckets** - If multipart uploads in a directory bucket are in progress, you can't delete the bucket until all the in-progress multipart uploads are aborted or completed.
- **Directory buckets** - For directory buckets, you must make requests for this API operation to the Regional endpoint. These endpoints support path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name`. Virtual-hosted-style requests aren't supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

Permissions

- **General purpose bucket permissions** - You must have the `s3:DeleteBucket` permission on the specified bucket in a policy.
- **Directory bucket permissions** - You must have the `s3express:DeleteBucket` permission in an IAM identity-based policy instead of a bucket policy. Cross-account access to this API operation isn't supported. This operation can only be performed by the AWS account that owns the resource. For more information about directory bucket policies and permissions, see [AWS Identity and Access Management \(IAM\) for S3 Express One Zone](#) in the *Amazon S3 User Guide*.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `s3express-control.region-code.amazonaws.com`.

The following operations are related to `DeleteBucket`:

- [CreateBucket](#)

- [DeleteObject](#)

⚠ Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
DELETE / HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

[Bucket](#)

Specifies the bucket being deleted.

Directory buckets - When you use this operation with a directory bucket, you must use path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name`. Virtual-hosted-style requests aren't supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must also follow the format `bucket-base-name--zone-id--x-s3` (for example, `DOC-EXAMPLE-BUCKET--usw2-az1--x-s3`). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*

Required: Yes

[x-amz-expected-bucket-owner](#)

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code `403 Forbidden` (access denied).

Note

For directory buckets, this header is not supported in this API operation. If you specify this header, the request fails with the HTTP status code 501 Not Implemented.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Examples

Sample Request for general purpose buckets

This request deletes the bucket named `amzn-s3-demo-bucket`.

```
DELETE / HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Wed, 01 Mar 2006 12:00:00 GMT
Authorization: authorization string
```

Sample Response for general purpose buckets

```
HTTP/1.1 204 No Content
x-amz-id-2: JuKZqmXuiwFeDQxhD7M8KtsKobSzWA1QEjLbTMTagkKdBX2z7I1/jGhDeJ3j6s80
x-amz-request-id: 32FE2CEB32F5EE25
Date: Wed, 01 Mar 2006 12:00:00 GMT
Connection: close
Server: AmazonS3
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteBucketAnalyticsConfiguration

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Deletes an analytics configuration for the bucket (specified by the analytics configuration ID).

To use this operation, you must have permissions to perform the `s3:PutAnalyticsConfiguration` action. The bucket owner has this permission by default. The bucket owner can grant this permission to others. For more information about permissions, see [Permissions Related to Bucket Subresource Operations](#) and [Managing Access Permissions to Your Amazon S3 Resources](#).

For information about the Amazon S3 analytics feature, see [Amazon S3 Analytics – Storage Class Analysis](#).

The following operations are related to `DeleteBucketAnalyticsConfiguration`:

- [GetBucketAnalyticsConfiguration](#)
- [ListBucketAnalyticsConfigurations](#)
- [PutBucketAnalyticsConfiguration](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
DELETE /?analytics&id=Id HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket from which an analytics configuration is deleted.

Required: Yes

id

The ID that identifies the analytics configuration.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Examples

Sample Request

The following DELETE request deletes the analytics configuration with the ID `list1`.

```
DELETE ?/analytics&id=list1 HTTP/1.1
```

```
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Wed, 14 May 2014 02:11:22 GMT
Authorization: signatureValue
```

Sample Response

The following successful response shows Amazon S3 returning a 204 No Content response. The analytics configuration with the ID list1 for the bucket has been removed.

```
HTTP/1.1 204 No Content
x-amz-id-2: 0FmFIWsh/
PpBuzZ0JFRC55ZGvmQW4SHJ7xVDqKwhEdJmf3q63RtrvH8ZuxW1Bo15
x-amz-request-id: 0CF038E9BCF63097
Date: Wed, 14 May 2014 02:11:22 GMT
Server: AmazonS3
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteBucketCors

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Deletes the cors configuration information set for the bucket.

To use this operation, you must have permission to perform the `s3:PutBucketCORS` action. The bucket owner has this permission by default and can grant this permission to others.

For information about cors, see [Enabling Cross-Origin Resource Sharing](#) in the *Amazon S3 User Guide*.

Related Resources

- [PutBucketCors](#)
- [RESTOPTIONSobject](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
DELETE /?cors HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

Specifies the bucket whose cors configuration is being deleted.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Examples

Sample Request

This example illustrates one usage of DeleteBucketCors.

```
DELETE /?cors HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Tue, 13 Dec 2011 19:14:42 GMT
Authorization: signatureValue
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteBucketEncryption

Service: Amazon S3

This implementation of the DELETE action resets the default encryption for the bucket as server-side encryption with Amazon S3 managed keys (SSE-S3).

Note

- **General purpose buckets** - For information about the bucket default encryption feature, see [Amazon S3 Bucket Default Encryption](#) in the *Amazon S3 User Guide*.
- **Directory buckets** - For directory buckets, there are only two supported options for server-side encryption: SSE-S3 and SSE-KMS. For information about the default encryption configuration in directory buckets, see [Setting default server-side encryption behavior for directory buckets](#).

Permissions

- **General purpose bucket permissions** - The `s3:PutEncryptionConfiguration` permission is required in a policy. The bucket owner has this permission by default. The bucket owner can grant this permission to others. For more information about permissions, see [Permissions Related to Bucket Operations](#) and [Managing Access Permissions to Your Amazon S3 Resources](#).
- **Directory bucket permissions** - To grant access to this API operation, you must have the `s3express:PutEncryptionConfiguration` permission in an IAM identity-based policy instead of a bucket policy. Cross-account access to this API operation isn't supported. This operation can only be performed by the AWS account that owns the resource. For more information about directory bucket policies and permissions, see [AWS Identity and Access Management \(IAM\) for S3 Express One Zone](#) in the *Amazon S3 User Guide*.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `s3express-control.region-code.amazonaws.com`.

The following operations are related to `DeleteBucketEncryption`:

- [PutBucketEncryption](#)

- [GetBucketEncryption](#)

⚠ Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
DELETE /?encryption HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket containing the server-side encryption configuration to delete.

Directory buckets - When you use this operation with a directory bucket, you must use path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name`. Virtual-hosted-style requests aren't supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must also follow the format `bucket-base-name--zone-id--x-s3` (for example, `DOC-EXAMPLE-BUCKET--usw2-az1--x-s3`). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Note

For directory buckets, this header is not supported in this API operation. If you specify this header, the request fails with the HTTP status code 501 Not Implemented.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Examples

Sample Request for a general purpose bucket

The following DELETE request resets the default encryption for the bucket as server-side encryption with Amazon S3 managed keys (SSE-S3).

```
DELETE ?/encryption HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Wed, 06 Sep 2017 12:00:00 GMT
Authorization: signatureValue
```

Sample Response for a general purpose bucket

The following successful response shows Amazon S3 returning a 204 No Content response confirming that default encryption for the bucket has been reset as server-side encryption with Amazon S3 managed keys (SSE-S3).

```
HTTP/1.1 204 No Content
x-amz-id-2: 0FmFIWsh/PpBuzZ0JFRC55ZGVmQW4SHJ7xVDqKwhEdJmf3q63RtrvH8ZuxW1Bo15
x-amz-request-id: 0CF038E9BCF63097
```

Date: Wed, 06 Sep 2017 12:00:00 GMT
Server: AmazonS3

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteBucketIntelligentTieringConfiguration

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Deletes the S3 Intelligent-Tiering configuration from the specified bucket.

The S3 Intelligent-Tiering storage class is designed to optimize storage costs by automatically moving data to the most cost-effective storage access tier, without performance impact or operational overhead. S3 Intelligent-Tiering delivers automatic cost savings in three low latency and high throughput access tiers. To get the lowest storage cost on data that can be accessed in minutes to hours, you can choose to activate additional archiving capabilities.

The S3 Intelligent-Tiering storage class is the ideal storage class for data with unknown, changing, or unpredictable access patterns, independent of object size or retention period. If the size of an object is less than 128 KB, it is not monitored and not eligible for auto-tiering. Smaller objects can be stored, but they are always charged at the Frequent Access tier rates in the S3 Intelligent-Tiering storage class.

For more information, see [Storage class for automatically optimizing frequently and infrequently accessed objects](#).

Operations related to DeleteBucketIntelligentTieringConfiguration include:

- [GetBucketIntelligentTieringConfiguration](#)
- [PutBucketIntelligentTieringConfiguration](#)
- [ListBucketIntelligentTieringConfigurations](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
DELETE /?intelligent-tiering&id=Id HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the Amazon S3 bucket whose configuration you want to modify or retrieve.

Required: Yes

id

The ID used to identify the S3 Intelligent-Tiering configuration.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteBucketInventoryConfiguration

Service: Amazon S3

Deletes an S3 Inventory configuration (identified by the inventory ID) from the bucket.

Note

Directory buckets - For directory buckets, you must make requests for this API operation to the Regional endpoint. These endpoints support path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name`. Virtual-hosted-style requests aren't supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

Permissions

To use this operation, you must have permissions to perform the `s3:PutInventoryConfiguration` action. The bucket owner has this permission by default. The bucket owner can grant this permission to others. For more information about permissions, see [Permissions Related to Bucket Subresource Operations](#) and [Managing Access Permissions to Your Amazon S3 Resources](#).

- **General purpose bucket permissions** - The `s3:PutInventoryConfiguration` permission is required in a policy. For more information about general purpose buckets permissions, see [Using Bucket Policies and User Policies](#) in the *Amazon S3 User Guide*.
- **Directory bucket permissions** - To grant access to this API operation, you must have the `s3express:PutInventoryConfiguration` permission in an IAM identity-based policy instead of a bucket policy. For more information about directory bucket policies and permissions, see [AWS Identity and Access Management \(IAM\) for S3 Express One Zone](#) in the *Amazon S3 User Guide*.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `s3express-control.region-code.amazonaws.com`.

For information about the Amazon S3 inventory feature, see [Amazon S3 Inventory](#).

Note

After deleting a configuration, Amazon S3 might still deliver one additional inventory report during a brief transition period while the system processes the deletion.

Operations related to `DeleteBucketInventoryConfiguration` include:

- [GetBucketInventoryConfiguration](#)
- [PutBucketInventoryConfiguration](#)
- [ListBucketInventoryConfigurations](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
DELETE /?inventory&id=Id HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket containing the inventory configuration to delete.

Directory buckets - When you use this operation with a directory bucket, you must use path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name` . Virtual-hosted-style requests aren't supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must also follow the format `bucket-base-name--zone-id--x-s3` (for

example, `DOC-EXAMPLE-BUCKET--usw2-az1--x-s3`). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*

Required: Yes

id

The ID used to identify the inventory configuration.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Note

For directory buckets, this header is not supported in this API operation. If you specify this header, the request fails with the HTTP status code 501 Not Implemented.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Examples

Sample Request

The following DELETE request deletes the inventory configuration with the ID `list1`.

```
DELETE ?/inventory&id=list1 HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Wed, 14 May 2014 02:11:22 GMT
Authorization: signatureValue
```

Sample Response

The following successful response shows Amazon S3 returning a 204 No Content response. The inventory configuration with the ID `list1` for the bucket has been removed.

```
HTTP/1.1 204 No Content
x-amz-id-2: 0FmFIWsh/PpBuzZ0JFRC55ZGVmQW4SHJ7xVDqKwhEdJmf3q63RtrvH8ZuxW1Bo15
x-amz-request-id: 0CF038E9BCF63097
Date: Wed, 14 May 2014 02:11:22 GMT
Server: AmazonS3
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteBucketLifecycle

Service: Amazon S3

Deletes the lifecycle configuration from the specified bucket. Amazon S3 removes all the lifecycle configuration rules in the lifecycle subresource associated with the bucket. Your objects never expire, and Amazon S3 no longer automatically deletes any objects on the basis of rules contained in the deleted lifecycle configuration.

Permissions

- **General purpose bucket permissions** - By default, all Amazon S3 resources are private, including buckets, objects, and related subresources (for example, lifecycle configuration and website configuration). Only the resource owner (that is, the AWS account that created it) can access the resource. The resource owner can optionally grant access permissions to others by writing an access policy. For this operation, a user must have the `s3:PutLifecycleConfiguration` permission.

For more information about permissions, see [Managing Access Permissions to Your Amazon S3 Resources](#).

- **Directory bucket permissions** - You must have the `s3express:PutLifecycleConfiguration` permission in an IAM identity-based policy to use this operation. Cross-account access to this API operation isn't supported. The resource owner can optionally grant access permissions to others by creating a role or user for them as long as they are within the same account as the owner and resource.

For more information about directory bucket policies and permissions, see [Authorizing Regional endpoint APIs with IAM](#) in the *Amazon S3 User Guide*.

Note

Directory buckets - For directory buckets, you must make requests for this API operation to the Regional endpoint. These endpoints support path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name`. Virtual-hosted-style requests aren't supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `s3express-control.region.amazonaws.com`.

For more information about the object expiration, see [Elements to Describe Lifecycle Actions](#).

Related actions include:

- [PutBucketLifecycleConfiguration](#)
- [GetBucketLifecycleConfiguration](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
DELETE /?lifecycle HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

[Bucket](#)

The bucket name of the lifecycle to delete.

Required: Yes

[x-amz-expected-bucket-owner](#)

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code `403 Forbidden` (access denied).

Note

This parameter applies to general purpose buckets only. It is not supported for directory bucket lifecycle configurations.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Examples

Sample Request

The following DELETE request deletes the lifecycle subresource from the specified general purpose bucket. This removes lifecycle configuration stored in the subresource.

```
DELETE /?lifecycle HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Wed, 14 Dec 2011 05:37:16 GMT
Authorization: signatureValue
```

Sample Response

The following successful response shows Amazon S3 returning a 204 No Content response. Objects in your general purpose bucket no longer expire.

```
HTTP/1.1 204 No Content
x-amz-id-2: Uuag1LuByRx9e6j50nimrSAMPLEtRPfTa0Aa==
x-amz-request-id: 656c76696e672SAMPLE5657374
```


Date: Wed, 14 Dec 2011 05:37:16 GMT
Connection: keep-alive
Server: AmazonS3

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteBucketMetadataConfiguration

Service: Amazon S3

Deletes an S3 Metadata configuration from a general purpose bucket. For more information, see [Accelerating data discovery with S3 Metadata](#) in the *Amazon S3 User Guide*.

Note

You can use the V2 DeleteBucketMetadataConfiguration API operation with V1 or V2 metadata configurations. However, if you try to use the V1 DeleteBucketMetadataTableConfiguration API operation with V2 configurations, you will receive an HTTP 405 Method Not Allowed error.

Permissions

To use this operation, you must have the `s3:DeleteBucketMetadataTableConfiguration` permission. For more information, see [Setting up permissions for configuring metadata tables](#) in the *Amazon S3 User Guide*.

Note

The IAM policy action name is the same for the V1 and V2 API operations.

The following operations are related to DeleteBucketMetadataConfiguration:

- [CreateBucketMetadataConfiguration](#)
- [GetBucketMetadataConfiguration](#)
- [UpdateBucketMetadataInventoryTableConfiguration](#)
- [UpdateBucketMetadataJournalTableConfiguration](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
DELETE /?metadataConfiguration HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The general purpose bucket that you want to remove the metadata configuration from.

Required: Yes

x-amz-expected-bucket-owner

The expected bucket owner of the general purpose bucket that you want to remove the metadata table configuration from.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)

- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteBucketMetadataTableConfiguration

Service: Amazon S3

Important

We recommend that you delete your S3 Metadata configurations by using the V2 [DeleteBucketMetadataTableConfiguration](#) API operation. We no longer recommend using the V1 DeleteBucketMetadataTableConfiguration API operation.

If you created your S3 Metadata configuration before July 15, 2025, we recommend that you delete and re-create your configuration by using [CreateBucketMetadataConfiguration](#) so that you can expire journal table records and create a live inventory table.

Deletes a V1 S3 Metadata configuration from a general purpose bucket. For more information, see [Accelerating data discovery with S3 Metadata](#) in the *Amazon S3 User Guide*.

Note

You can use the V2 DeleteBucketMetadataConfiguration API operation with V1 or V2 metadata table configurations. However, if you try to use the V1 DeleteBucketMetadataTableConfiguration API operation with V2 configurations, you will receive an HTTP 405 Method Not Allowed error.

Make sure that you update your processes to use the new V2 API operations (CreateBucketMetadataConfiguration, GetBucketMetadataConfiguration, and DeleteBucketMetadataConfiguration) instead of the V1 API operations.

Permissions

To use this operation, you must have the `s3:DeleteBucketMetadataTableConfiguration` permission. For more information, see [Setting up permissions for configuring metadata tables](#) in the *Amazon S3 User Guide*.

The following operations are related to DeleteBucketMetadataTableConfiguration:

- [CreateBucketMetadataTableConfiguration](#)
- [GetBucketMetadataTableConfiguration](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
DELETE /?metadataTable HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The general purpose bucket that you want to remove the metadata table configuration from.

Required: Yes

x-amz-expected-bucket-owner

The expected bucket owner of the general purpose bucket that you want to remove the metadata table configuration from.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteBucketMetricsConfiguration

Service: Amazon S3

Deletes a metrics configuration for the Amazon CloudWatch request metrics (specified by the metrics configuration ID) from the bucket. Note that this doesn't include the daily storage metrics.

Note

Directory buckets - For directory buckets, you must make requests for this API operation to the Regional endpoint. These endpoints support path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name` . Virtual-hosted-style requests aren't supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

Permissions

To use this operation, you must have permissions to perform the `s3:PutMetricsConfiguration` action. The bucket owner has this permission by default. The bucket owner can grant this permission to others. For more information about permissions, see [Permissions Related to Bucket Subresource Operations](#) and [Managing Access Permissions to Your Amazon S3 Resources](#).

- **General purpose bucket permissions** - The `s3:PutMetricsConfiguration` permission is required in a policy. For more information about general purpose buckets permissions, see [Using Bucket Policies and User Policies](#) in the *Amazon S3 User Guide*.
- **Directory bucket permissions** - To grant access to this API operation, you must have the `s3express:PutMetricsConfiguration` permission in an IAM identity-based policy instead of a bucket policy. Cross-account access to this API operation isn't supported. This operation can only be performed by the AWS account that owns the resource. For more information about directory bucket policies and permissions, see [AWS Identity and Access Management \(IAM\) for S3 Express One Zone](#) in the *Amazon S3 User Guide*.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `s3express-control.region-code.amazonaws.com`.

For information about CloudWatch request metrics for Amazon S3, see [Monitoring Metrics with Amazon CloudWatch](#).

The following operations are related to `DeleteBucketMetricsConfiguration`:

- [GetBucketMetricsConfiguration](#)
- [PutBucketMetricsConfiguration](#)
- [ListBucketMetricsConfigurations](#)
- [Monitoring Metrics with Amazon CloudWatch](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
DELETE /?metrics&id=Id HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

[Bucket](#)

The name of the bucket containing the metrics configuration to delete.

Directory buckets - When you use this operation with a directory bucket, you must use path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name`. Virtual-hosted-style requests aren't supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must also follow the format `bucket-base-name--zone-id--x-s3` (for example, `DOC-EXAMPLE-BUCKET--usw2-az1--x-s3`). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*.

Required: Yes

id

The ID used to identify the metrics configuration. The ID has a 64 character limit and can only contain letters, numbers, periods, dashes, and underscores.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Note

For directory buckets, this header is not supported in this API operation. If you specify this header, the request fails with the HTTP status code 501 Not Implemented.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Examples

Sample Request

Delete the metric configuration with a specified ID, which disables the CloudWatch metrics with the ExampleMetrics value for the FilterId dimension.

```
DELETE /?metrics&id=ExampleMetrics HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
x-amz-date: Thu, 15 Nov 2016 00:17:21 GMT
Authorization: signatureValue
```

Sample Response

Delete the metric configuration with a specified ID, which disables the CloudWatch metrics with the ExampleMetrics value for the FilterId dimension.

```
HTTP/1.1 204 No Content
x-amz-id-2:
ITnGT1y4REXAMPLEPi4hk1TXouTf0hccUjo0iCPEXAMPLEutBj3M7fPGLW02SEWp
x-amz-request-id: 51991EXAMPLE5321
Date: Thu, 15 Nov 2016 00:17:22 GMT
Server: AmazonS3
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteBucketOwnershipControls

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Removes `OwnershipControls` for an Amazon S3 bucket. To use this operation, you must have the `s3:PutBucketOwnershipControls` permission. For more information about Amazon S3 permissions, see [Specifying Permissions in a Policy](#).

For information about Amazon S3 Object Ownership, see [Using Object Ownership](#).

The following operations are related to `DeleteBucketOwnershipControls`:

- [GetBucketOwnershipControls](#)
- [PutBucketOwnershipControls](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
DELETE /?ownershipControls HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The Amazon S3 bucket whose `OwnershipControls` you want to delete.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Examples

Sample DeleteBucketOwnershipControls Request

This example illustrates one usage of DeleteBucketOwnershipControls.

```
DELETE /example-bucket?/ownershipControls HTTP/1.1
Host: examplebucket.s3.<Region>.amazonaws.com
Date: Thu, 18 Jun 2017 00:17:22 GMT
Authorization: signatureValue;
```

Sample DeleteBucketOwnershipControls Response

This example illustrates one usage of DeleteBucketOwnershipControls.

```
HTTP/1.1 204 No Content
x-amz-id-2: dVrxJD3XHDcjZHFtd7eSB+ovpY8hQ6kSe9jPzyRVkWp27cij05qV1pTIvz/
hjlsrupiy9gEkSdw=
```

```
x-amz-request-id: 4BFC0B777B448C97  
Date: Thu, 18 Jun 2020 22:54:03 GMT  
Server: AmazonS3
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteBucketPolicy

Service: Amazon S3

Deletes the policy of a specified bucket.

Note

Directory buckets - For directory buckets, you must make requests for this API operation to the Regional endpoint. These endpoints support path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name` . Virtual-hosted-style requests aren't supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

Permissions

If you are using an identity other than the root user of the AWS account that owns the bucket, the calling identity must both have the DeleteBucketPolicy permissions on the specified bucket and belong to the bucket owner's account in order to use this operation.

If you don't have DeleteBucketPolicy permissions, Amazon S3 returns a 403 Access Denied error. If you have the correct permissions, but you're not using an identity that belongs to the bucket owner's account, Amazon S3 returns a 405 Method Not Allowed error.

Important

To ensure that bucket owners don't inadvertently lock themselves out of their own buckets, the root principal in a bucket owner's AWS account can perform the GetBucketPolicy, PutBucketPolicy, and DeleteBucketPolicy API actions, even if their bucket policy explicitly denies the root principal's access. Bucket owner root principals can only be blocked from performing these API actions by VPC endpoint policies and AWS Organizations policies.

- **General purpose bucket permissions** - The `s3:DeleteBucketPolicy` permission is required in a policy. For more information about general purpose buckets bucket policies, see [Using Bucket Policies and User Policies](#) in the *Amazon S3 User Guide*.

- **Directory bucket permissions** - To grant access to this API operation, you must have the `s3express:DeleteBucketPolicy` permission in an IAM identity-based policy instead of a bucket policy. Cross-account access to this API operation isn't supported. This operation can only be performed by the AWS account that owns the resource. For more information about directory bucket policies and permissions, see [AWS Identity and Access Management \(IAM\) for S3 Express One Zone](#) in the *Amazon S3 User Guide*.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `s3express-control.region-code.amazonaws.com`.

The following operations are related to `DeleteBucketPolicy`

- [CreateBucket](#)
- [DeleteObject](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
DELETE /?policy HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

[Bucket](#)

The bucket name.

Directory buckets - When you use this operation with a directory bucket, you must use path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name` . Virtual-hosted-style requests aren't supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must also follow the format `bucket-base-name--zone-id--x-s3` (for example, `DOC-EXAMPLE-BUCKET--usw2-az1--x-s3`). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*

Required: Yes

[x-amz-expected-bucket-owner](#)

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Note

For directory buckets, this header is not supported in this API operation. If you specify this header, the request fails with the HTTP status code 501 Not Implemented.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Examples

Sample Request for general purpose buckets

This request deletes the bucket named `amzn-s3-demo-bucket`.

```
DELETE /?policy HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Tue, 04 Apr 2010 20:34:56 GMT
Authorization: signatureValue
```

Sample Response for general purpose buckets

This example illustrates one usage of DeleteBucketPolicy.

```
HTTP/1.1 204 No Content
x-amz-id-2: Uuag1LuByRx9e6j50nimrSAMPLEtRPfTa0Fg==
x-amz-request-id: 656c76696e672SAMPLE5657374
Date: Tue, 04 Apr 2010 20:34:56 GMT
Connection: keep-alive
Server: AmazonS3
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteBucketReplication

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Deletes the replication configuration from the bucket.

To use this operation, you must have permissions to perform the `s3:PutReplicationConfiguration` action. The bucket owner has these permissions by default and can grant it to others. For more information about permissions, see [Permissions Related to Bucket Subresource Operations](#) and [Managing Access Permissions to Your Amazon S3 Resources](#).

Note

It can take a while for the deletion of a replication configuration to fully propagate.

For information about replication configuration, see [Replication](#) in the *Amazon S3 User Guide*.

The following operations are related to `DeleteBucketReplication`:

- [PutBucketReplication](#)
- [GetBucketReplication](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
DELETE /?replication HTTP/1.1
Host: Bucket.s3.amazonaws.com
```

x-amz-expected-bucket-owner: *ExpectedBucketOwner*

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

HTTP/1.1 204

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Examples

Sample Request

The following DELETE request deletes the replication subresource from the specified bucket. This removes the replication configuration that is set for the bucket.

```
DELETE /?replication HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Wed, 11 Feb 2015 05:37:16 GMT
```

```
20150211T171320Z
```

```
Authorization: authorization string
```

Sample Response

When the replication subresource has been deleted, Amazon S3 returns a 204 No Content response. It will not replicate new objects that are stored in the examplebucket bucket.

```
HTTP/1.1 204 No Content
x-amz-id-2: Uuag1LuByRx9e6j50nimrSAMPLEtRPfTa0Aa==
x-amz-request-id: 656c76696e672example
Date: Wed, 11 Feb 2015 05:37:16 GMT
Connection: keep-alive
Server: AmazonS3
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteBucketTagging

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Deletes tags from the general purpose bucket if attribute based access control (ABAC) is not enabled for the bucket. When you [enable ABAC for a general purpose bucket](#), you can no longer use this operation for that bucket and must use [UntagResource](#) instead.

To use this operation, you must have permission to perform the `s3:PutBucketTagging` action. By default, the bucket owner has this permission and can grant this permission to others.

The following operations are related to DeleteBucketTagging:

- [GetBucketTagging](#)
- [PutBucketTagging](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
DELETE /?tagging HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket that has the tag set to be removed.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Examples

Sample Request

The following DELETE request deletes the tag set from the specified bucket.

```
DELETE /?tagging HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Wed, 14 Dec 2011 05:37:16 GMT
Authorization: signatureValue
```

Sample Response

The following successful response shows Amazon S3 returning a 204 No Content response. The tag set for the bucket has been removed.

```
HTTP/1.1 204 No Content
Date: Wed, 25 Nov 2009 12:00:00 GMT
Connection: close
Server: AmazonS3
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteBucketWebsite

Service: Amazon S3

Note

This operation is not supported for directory buckets.

This action removes the website configuration for a bucket. Amazon S3 returns a 200 OK response upon successfully deleting a website configuration on the specified bucket. You will get a 200 OK response if the website configuration you are trying to delete does not exist on the bucket. Amazon S3 returns a 404 response if the bucket specified in the request does not exist.

This DELETE action requires the `S3:DeleteBucketWebsite` permission. By default, only the bucket owner can delete the website configuration attached to a bucket. However, bucket owners can grant other users permission to delete the website configuration by writing a bucket policy granting them the `S3:DeleteBucketWebsite` permission.

For more information about hosting websites, see [Hosting Websites on Amazon S3](#).

The following operations are related to DeleteBucketWebsite:

- [GetBucketWebsite](#)
- [PutBucketWebsite](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
DELETE /?website HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name for which you want to remove the website configuration.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Examples

Sample Request

This request deletes the website configuration on the specified bucket.

```
DELETE ?website HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Thu, 27 Jan 2011 12:00:00 GMT
Authorization: signatureValue
```

Sample Response

This example illustrates one usage of DeleteBucketWebsite.

```
HTTP/1.1 204 No Content
x-amz-id-2: aws-s3integ-s3ws-31008.sea31.amazon.com
x-amz-request-id: AF1DD829D3B49707
Date: Thu, 03 Feb 2011 22:10:26 GMT
Server: AmazonS3
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteObject

Service: Amazon S3

Removes an object from a bucket. The behavior depends on the bucket's versioning state:

- If bucket versioning is not enabled, the operation permanently deletes the object.
- If bucket versioning is enabled, the operation inserts a delete marker, which becomes the current version of the object. To permanently delete an object in a versioned bucket, you must include the object's `versionId` in the request. For more information about versioning-enabled buckets, see [Deleting object versions from a versioning-enabled bucket](#).
- If bucket versioning is suspended, the operation removes the object that has a null `versionId`, if there is one, and inserts a delete marker that becomes the current version of the object. If there isn't an object with a null `versionId`, and all versions of the object have a `versionId`, Amazon S3 does not remove the object and only inserts a delete marker. To permanently delete an object that has a `versionId`, you must include the object's `versionId` in the request. For more information about versioning-suspended buckets, see [Deleting objects from versioning-suspended buckets](#).

Note

- **Directory buckets** - S3 Versioning isn't enabled and supported for directory buckets. For this API operation, only the null value of the version ID is supported by directory buckets. You can only specify null to the `versionId` query parameter in the request.
- **Directory buckets** - For directory buckets, you must make requests for this API operation to the Zonal endpoint. These endpoints support virtual-hosted-style requests in the format `https://amzn-s3-demo-bucket.s3express-zone-id.region-code.amazonaws.com/key-name`. Path-style requests are not supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

To remove a specific version, you must use the `versionId` query parameter. Using this query parameter permanently deletes the version. If the object deleted is a delete marker, Amazon S3 sets the response header `x-amz-delete-marker` to true.

If the object you want to delete is in a bucket where the bucket versioning configuration is MFA Delete enabled, you must include the `x-amz-mfa` request header in the `DELETE` `versionId` request. Requests that include `x-amz-mfa` must use HTTPS. For more information about MFA Delete, see [Using MFA Delete](#) in the *Amazon S3 User Guide*. To see sample requests that use versioning, see [Sample Request](#).

 Note

Directory buckets - MFA delete is not supported by directory buckets.

You can delete objects by explicitly calling `DELETE` Object or calling ([PutBucketLifecycle](#)) to enable Amazon S3 to remove them for you. If you want to block users or accounts from removing or deleting objects from your bucket, you must deny them the `s3:DeleteObject`, `s3:DeleteObjectVersion`, and `s3:PutLifecycleConfiguration` actions.

 Note

Directory buckets - S3 Lifecycle is not supported by directory buckets.

Permissions

- **General purpose bucket permissions** - The following permissions are required in your policies when your `DeleteObjects` request includes specific headers.
- **`s3:DeleteObject`** - To delete an object from a bucket, you must always have the `s3:DeleteObject` permission.
- **`s3:DeleteObjectVersion`** - To delete a specific version of an object from a versioning-enabled bucket, you must have the `s3:DeleteObjectVersion` permission.

 Note

If the `s3:DeleteObject` or `s3:DeleteObjectVersion` permissions are explicitly denied in your bucket policy, attempts to delete any unversioned objects result in a `403 Access Denied` error.

- **Directory bucket permissions** - To grant access to this API operation on a directory bucket, we recommend that you use the [CreateSession](#) API operation for session-based

authorization. Specifically, you grant the `s3express:CreateSession` permission to the directory bucket in a bucket policy or an IAM identity-based policy. Then, you make the `CreateSession` API call on the bucket to obtain a session token. With the session token in your request header, you can make API requests to this operation. After the session token expires, you make another `CreateSession` API call to generate a new session token for use. AWS CLI or SDKs create session and refresh the session token automatically to avoid service interruptions when a session expires. For more information about authorization, see [CreateSession](#).

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `Bucket-name.s3express-zone-id.region-code.amazonaws.com`.

The following action is related to `DeleteObject`:

- [PutObject](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Note

The `If-Match` header is supported for both general purpose and directory buckets. `IfMatchLastModifiedTime` and `IfMatchSize` is only supported for directory buckets.

Request Syntax

```
DELETE /Key+?versionId=VersionId HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-mfa: MFA
x-amz-request-payer: RequestPayer
x-amz-bypass-governance-retention: BypassGovernanceRetention
```

```
x-amz-expected-bucket-owner: ExpectedBucketOwner  
If-Match: IfMatch  
x-amz-if-match-last-modified-time: IfMatchLastModifiedTime  
x-amz-if-match-size: IfMatchSize
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name of the bucket containing the object.

Directory buckets - When you use this operation with a directory bucket, you must use virtual-hosted-style requests in the format `Bucket-name.s3express-zone-id.region-code.amazonaws.com`. Path-style requests are not supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must follow the format `bucket-base-name--zone-id--x-s3` (for example, `amzn-s3-demo-bucket--usw2-az1--x-s3`). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form `AccessPointName-AccountId.s3-accesspoint.Region.amazonaws.com`. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

Note

Object Lambda access points are not supported by directory buckets.

S3 on Outposts - When you use this action with S3 on Outposts, you must direct requests to the S3 on Outposts hostname. The S3 on Outposts hostname takes the form `AccessPointName-AccountId.outpostID.s3-outposts.Region.amazonaws.com`. When you use this action with S3 on Outposts, the destination bucket must be the Outposts

access point ARN or the access point alias. For more information about S3 on Outposts, see [What is S3 on Outposts?](#) in the *Amazon S3 User Guide*.

Required: Yes

If-Match

Deletes the object if the ETag (entity tag) value provided during the delete operation matches the ETag of the object in S3. If the ETag values do not match, the operation returns a 412 Precondition Failed error.

Expects the ETag value as a string. If-Match does accept a string value of an '*' (asterisk) character to denote a match of any ETag.

For more information about conditional requests, see [RFC 7232](#).

Key

Key name of the object to delete.

Length Constraints: Minimum length of 1.

Required: Yes

versionId

Version ID used to reference a specific version of the object.

Note

For directory buckets in this API operation, only the null value of the version ID is supported.

x-amz-bypass-governance-retention

Indicates whether S3 Object Lock should bypass Governance-mode restrictions to process this operation. To use this header, you must have the `s3:BypassGovernanceRetention` permission.

Note

This functionality is not supported for directory buckets.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-if-match-last-modified-time

If present, the object is deleted only if its modification times matches the provided Timestamp. If the Timestamp values do not match, the operation returns a 412 Precondition Failed error. If the Timestamp matches or if the object doesn't exist, the operation returns a 204 Success (No Content) response.

Note

This functionality is only supported for directory buckets.

x-amz-if-match-size

If present, the object is deleted only if its size matches the provided size in bytes. If the Size value does not match, the operation returns a 412 Precondition Failed error. If the Size matches or if the object doesn't exist, the operation returns a 204 Success (No Content) response.

Note

This functionality is only supported for directory buckets.

Important

You can use the If-Match, x-amz-if-match-last-modified-time and x-amz-if-match-size conditional headers in conjunction with each-other or individually.

x-amz-mfa

The concatenation of the authentication device's serial number, a space, and the value that is displayed on your authentication device. Required to permanently delete a versioned object if versioning is configured with MFA delete enabled.

Note

This functionality is not supported for directory buckets.

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
x-amz-delete-marker: DeleteMarker
x-amz-version-id: VersionId
x-amz-request-charged: RequestCharged
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response.

The response returns the following HTTP headers.

x-amz-delete-marker

Indicates whether the specified object version that was permanently deleted was (true) or was not (false) a delete marker before deletion. In a simple DELETE, this header indicates whether (true) or not (false) the current version of the object is a delete marker. To learn more about delete markers, see [Working with delete markers](#).

Note

This functionality is not supported for directory buckets.

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

x-amz-version-id

Returns the version ID of the delete marker created as a result of the DELETE operation.

Note

This functionality is not supported for directory buckets.

Examples

Sample Request for general purpose buckets

The following request deletes the object `my-second-image.jpg`.

```
DELETE /my-second-image.jpg HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Wed, 12 Oct 2009 17:50:00 GMT
Authorization: authorization string
Content-Type: text/plain
```

Sample Response for general purpose buckets

This example illustrates one usage of DeleteObject.

```
HTTP/1.1 204 NoContent
x-amz-id-2: LriYPLdm0dAiIfgSm/F1YsViT1LW94/xUQxMsF7xiEb1a0wiI0Ix1+zbwZ163pt7
x-amz-request-id: 0A49CE4060975EAC
Date: Wed, 12 Oct 2009 17:50:00 GMT
Content-Length: 0
Connection: close
Server: AmazonS3
```

Sample Request for general purpose buckets: Deleting a specified version of an object

The following request deletes the specified version of the object `my-third-image.jpg`.

```
DELETE /my-third-image.jpg?
versionId=UI0RUNfndfiufdisojhr398493jfdkjFJjkndnqUifhnw89493jJFJ HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Wed, 12 Oct 2009 17:50:00 GMT
Authorization: authorization string
Content-Type: text/plain
Content-Length: 0
```

Sample Response for general purpose buckets

This example illustrates one usage of DeleteObject.

```
HTTP/1.1 204 NoContent
```

```
x-amz-id-2: LriYPLdm0dAiIfgSm/F1YsViT1LW94/xUQxMsF7xiEb1a0wiI0Ix1+zbwZ163pt7
x-amz-request-id: 0A49CE4060975EAC
x-amz-version-id: UIORUnfndfiufdisojhr398493jfdkjFJjkndnqUifhnw89493jJFJ
Date: Wed, 12 Oct 2009 17:50:00 GMT
Content-Length: 0
Connection: close
Server: AmazonS3
```

Sample Response for general purpose buckets: If the object deleted is a delete marker

This example illustrates one usage of DeleteObject.

```
HTTP/1.1 204 NoContent
x-amz-id-2: LriYPLdm0dAiIfgSm/F1YsViT1LW94/xUQxMsF7xiEb1a0wiI0Ix1
+zbwZ163pt7
x-amz-request-id: 0A49CE4060975EAC
x-amz-version-id: 3/L4kqtJlcpXroDTDmJ+rmSpXd3dIbrHY
+MTRCxf3vjVBH40Nr8X8gdRQBpUMLUo
x-amz-delete-marker: true
Date: Wed, 12 Oct 2009 17:50:00 GMT
Content-Length: 0
Connection: close
Server: AmazonS3
```

Sample Request for general purpose buckets: Deleting a specified version of an object in an MFA-enabled bucket

The following request deletes the specified version of the object `my-third-image.jpg`, which is stored in an MFA-enabled bucket.

```
DELETE /my-third-image.jpg?versionId=UIORUnfndfiuf HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Wed, 12 Oct 2009 17:50:00 GMT
x-amz-mfa:[SerialNumber] [AuthenticationCode]
Authorization: authorization string
Content-Type: text/plain
Content-Length: 0
```

Sample Response for general purpose buckets

This example illustrates one usage of DeleteObject.

```
HTTP/1.1 204 NoContent
x-amz-id-2: LriYPLdm0dAiIfgSm/F1YsViT1LW94/xUQxMsF7xiEb1a0wiIOIx1
+zbwZ163pt7
x-amz-request-id: 0A49CE4060975EAC
x-amz-version-id: UIORUnfndfiuf
Date: Wed, 12 Oct 2009 17:50:00 GMT
Content-Length: 0
Connection: close
Server: AmazonS3
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteObjectAnnotation

Service: Amazon S3

Deletes a specific annotation from an Amazon S3 object. Use the `x-amz-object-if-match` header to perform a conditional delete that only succeeds if the object's ETag matches the provided value, preventing race conditions during concurrent updates.

Deleting an annotation is permanent. Annotations are not independently versioned, so there is no delete marker or way to recover a deleted annotation.

To use this operation, you must have the `s3:DeleteObjectAnnotation` permission. If the object is protected by Object Lock in governance mode, you must also include the `x-amz-bypass-governance-retention` header.

Note

Annotations are not supported by the following features: S3 Inventory Reports, API Gateway, S3 Storage Lens, Amazon S3 File Gateway, Amazon FSx, S3 on Outposts, and S3 Express One Zone (directory buckets).

The following operations are related to `DeleteObjectAnnotation`:

- [PutObjectAnnotation](#)
- [GetObjectAnnotation](#)
- [ListObjectAnnotations](#)

Request Syntax

```
DELETE /{Key+}?annotation&annotationName=AnnotationName&versionId=VersionId HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-request-payer: RequestPayer
x-amz-expected-bucket-owner: ExpectedBucketOwner
x-amz-object-if-match: ObjectIfMatch
```

URI Request Parameters

The request uses the following URI parameters.

annotationName

The name of the annotation to delete. Annotation names are UTF-8 encoded and cannot start with `aws` or `s3` (case-insensitive).

Length Constraints: Minimum length of 1. Maximum length of 512 bytes.

Required: Yes

Bucket

The name of the bucket that contains the object.

Required: Yes

Key

The object key.

Length Constraints: Minimum length of 1.

Required: Yes

versionId

The version ID of the object.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner.

x-amz-object-if-match

If specified, the operation only succeeds if the object's ETag matches the provided value.

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
x-amz-object-version-id: ObjectVersionId
x-amz-request-charged: RequestCharged
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response.

The response returns the following HTTP headers.

x-amz-object-version-id

The version ID of the object that the annotation was deleted from.

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

Errors

NoSuchBucket

The specified bucket does not exist.

HTTP Status Code: 404

NoSuchKey

The specified key does not exist.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteObjects

Service: Amazon S3

This operation enables you to delete multiple objects from a bucket using a single HTTP request. If you know the object keys that you want to delete, then this operation provides a suitable alternative to sending individual delete requests, reducing per-request overhead.

The request can contain a list of up to 1,000 keys that you want to delete. In the XML, you provide the object key names, and optionally, version IDs if you want to delete a specific version of the object from a versioning-enabled bucket. For each key, Amazon S3 performs a delete operation and returns the result of that delete, success or failure, in the response. If the object specified in the request isn't found, Amazon S3 confirms the deletion by returning the result as deleted.

Note

- **Directory buckets** - S3 Versioning isn't enabled and supported for directory buckets.
- **Directory buckets** - For directory buckets, you must make requests for this API operation to the Zonal endpoint. These endpoints support virtual-hosted-style requests in the format `https://amzn-s3-demo-bucket.s3express-zone-id.region-code.amazonaws.com/key-name` . Path-style requests are not supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

The operation supports two modes for the response: verbose and quiet. By default, the operation uses verbose mode in which the response includes the result of deletion of each key in your request. In quiet mode the response includes only keys where the delete operation encountered an error. For a successful deletion in a quiet mode, the operation does not return any information about the delete in the response body.

When performing this action on an MFA Delete enabled bucket, that attempts to delete any versioned objects, you must include an MFA token. If you do not provide one, the entire request will fail, even if there are non-versioned objects you are trying to delete. If you provide an invalid token, whether there are versioned keys in the request or not, the entire Multi-Object Delete request will fail. For information about MFA Delete, see [MFA Delete](#) in the *Amazon S3 User Guide*.

Note

Directory buckets - MFA delete is not supported by directory buckets.

Permissions

- **General purpose bucket permissions** - The following permissions are required in your policies when your `DeleteObjects` request includes specific headers.
- **s3:DeleteObject** - To delete an object from a bucket, you must always specify the `s3:DeleteObject` permission.
- **s3:DeleteObjectVersion** - To delete a specific version of an object from a versioning-enabled bucket, you must specify the `s3:DeleteObjectVersion` permission.

Note

If the `s3:DeleteObject` or `s3:DeleteObjectVersion` permissions are explicitly denied in your bucket policy, attempts to delete any unversioned objects result in a `403 Access Denied` error.

- **Directory bucket permissions** - To grant access to this API operation on a directory bucket, we recommend that you use the [CreateSession](#) API operation for session-based authorization. Specifically, you grant the `s3express:CreateSession` permission to the directory bucket in a bucket policy or an IAM identity-based policy. Then, you make the `CreateSession` API call on the bucket to obtain a session token. With the session token in your request header, you can make API requests to this operation. After the session token expires, you make another `CreateSession` API call to generate a new session token for use. AWS CLI or SDKs create session and refresh the session token automatically to avoid service interruptions when a session expires. For more information about authorization, see [CreateSession](#).

Content-MD5 request header

- **General purpose bucket** - The Content-MD5 request header is required for all Multi-Object Delete requests. Amazon S3 uses the header value to ensure that your request body has not been altered in transit.
- **Directory bucket** - The Content-MD5 request header or a additional checksum request header (including `x-amz-checksum-crc32`, `x-amz-checksum-crc32c`, `x-amz-`

checksum-sha1, or x-amz-checksum-sha256) is required for all Multi-Object Delete requests.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is *Bucket-name.s3express-zone-id.region-code.amazonaws.com*.

The following operations are related to DeleteObjects:

- [CreateMultipartUpload](#)
- [UploadPart](#)
- [CompleteMultipartUpload](#)
- [ListParts](#)
- [AbortMultipartUpload](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is my file.txt, containing two spaces after my, you must URL encode this value to my%20%20file.txt.

Request Syntax

```
POST /?delete HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-mfa: MFA
x-amz-request-payer: RequestPayer
x-amz-bypass-governance-retention: BypassGovernanceRetention
x-amz-expected-bucket-owner: ExpectedBucketOwner
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
<?xml version="1.0" encoding="UTF-8"?>
<Delete xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Object>
    <ETag>string</ETag>
    <Key>string</Key>
    <LastModifiedTime>timestamp</LastModifiedTime>
    <Size>long</Size>
```

```
<VersionId>string</VersionId>
</Object>
...
<Quiet>boolean</Quiet>
</Delete>
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name containing the objects to delete.

Directory buckets - When you use this operation with a directory bucket, you must use virtual-hosted-style requests in the format *Bucket-name.s3express-zone-id.region-code.amazonaws.com*. Path-style requests are not supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must follow the format *bucket-base-name--zone-id--x-s3* (for example, *amzn-s3-demo-bucket--usw2-az1--x-s3*). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form *AccessPointName-AccountId.s3-accesspoint.Region.amazonaws.com*. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

Note

Object Lambda access points are not supported by directory buckets.

S3 on Outposts - When you use this action with S3 on Outposts, you must direct requests to the S3 on Outposts hostname. The S3 on Outposts hostname takes the form

`AccessPointName-AccountId.outpostID.s3-outposts.Region.amazonaws.com`. When you use this action with S3 on Outposts, the destination bucket must be the Outposts access point ARN or the access point alias. For more information about S3 on Outposts, see [What is S3 on Outposts?](#) in the *Amazon S3 User Guide*.

Required: Yes

x-amz-bypass-governance-retention

Specifies whether you want to delete this object even if it has a Governance-type Object Lock in place. To use this header, you must have the `s3:BypassGovernanceRetention` permission.

Note

This functionality is not supported for directory buckets.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-mfa

The concatenation of the authentication device's serial number, a space, and the value that is displayed on your authentication device. Required to permanently delete a versioned object if versioning is configured with MFA delete enabled.

When performing the `DeleteObjects` operation on an MFA delete enabled bucket, which attempts to delete the specified versioned objects, you must include an MFA token. If you don't provide an MFA token, the entire request will fail, even if there are non-versioned objects that you are trying to delete. If you provide an invalid token, whether there are versioned object keys in the request or not, the entire Multi-Object Delete request will fail. For information about MFA Delete, see [MFA Delete](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

x-amz-sdk-checksum-algorithm

Indicates the algorithm used to create the checksum for the object when you use the SDK. This header will not provide any additional functionality if you don't use the SDK. When you send this header, there must be a corresponding `x-amz-checksum-algorithm` or `x-amz-trailer` header sent. Otherwise, Amazon S3 fails the request with the HTTP status code 400 Bad Request.

For the `x-amz-checksum-algorithm` header, replace *algorithm* with the supported algorithm from the following list:

- CRC32
- CRC32C
- CRC64NVME
- MD5
- SHA1
- SHA256
- SHA512
- XXHASH3
- XXHASH64
- XXHASH128

For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

If the individual checksum value you provide through `x-amz-checksum-algorithm` doesn't match the checksum algorithm you set through `x-amz-sdk-checksum-algorithm`, Amazon S3 fails the request with a `BadDigest` error.

If you provide an individual checksum, Amazon S3 ignores any provided `ChecksumAlgorithm` parameter.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 | XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in XML format.

Delete

Root level tag for the Delete parameters.

Required: Yes

Object

The object to delete.

Note

Directory buckets - For directory buckets, an object that's composed entirely of whitespace characters is not supported by the `DeleteObjects` API operation. The request will receive a `400 Bad Request` error and none of the objects in the request will be deleted.

Type: Array of [ObjectIdentifier](#) data types

Required: Yes

Quiet

Element to enable quiet mode for the request. When you add this element, you must set its value to `true`.

Type: Boolean

Required: No

Response Syntax

```
HTTP/1.1 200
x-amz-request-charged: RequestCharged
<?xml version="1.0" encoding="UTF-8"?>
<DeleteResult>
  <Deleted>
    <DeleteMarker>boolean</DeleteMarker>
    <DeleteMarkerVersionId>string</DeleteMarkerVersionId>
    <Key>string</Key>
    <VersionId>string</VersionId>
  </Deleted>
  ...
  <Error>
    <Code>string</Code>
    <Key>string</Key>
    <Message>string</Message>
    <VersionId>string</VersionId>
  </Error>
  ...
</DeleteResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

The following data is returned in XML format by the service.

DeleteResult

Root level tag for the DeleteResult parameters.

Required: Yes

Deleted

Container element for a successful delete. It identifies the object that was successfully deleted.

Type: Array of [DeletedObject](#) data types

Error

Container for a failed delete action that describes the object that Amazon S3 attempted to delete and the error it encountered.

Type: Array of [Error](#) data types

Examples

Sample Request for general purpose buckets: Multi-object delete resulting in mixed success/error response

This example illustrates a Multi-Object Delete request to delete objects that result in mixed success and errors response. The following request deletes two objects from a bucket (amzn-s3-demo-bucket). In this example, the requester does not have permission to delete the sample2.txt object.

```
POST /?delete HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Accept: */*
x-amz-date: Wed, 30 Nov 2011 03:39:05 GMT
Content-MD5: p5/WA/oEr30qrEE121PAqw==
Authorization: AWS AKIAIOSFODNN7EXAMPLE:W0qPYCL6JwkZAD1ei6hp9XZiIee=
Content-Length: 125
Connection: Keep-Alive
```

```
<Delete>
<Object>
<Key>sample1.txt</Key>
</Object>
<Object>
<Key>sample2.txt</Key>
</Object>
</Delete>
```

Sample Response for general purpose buckets

The response includes a `DeleteResult` element that includes a `Deleted` element for the item that Amazon S3 successfully deleted and an `Error` element that Amazon S3 did not delete because you didn't have permission to delete the object.

```
HTTP/1.1 200 OK
x-amz-id-2: 5h4FxSNCUS7wP5z92eGCWDshNpMnRuXvETa4HH3LvVH6VAIr0jU7tH9kM7X+njXx
x-amz-request-id: A437B3B641629AEE
Date: Fri, 02 Dec 2011 01:53:42 GMT
Content-Type: application/xml
Server: AmazonS3
Content-Length: 251

<?xml version="1.0" encoding="UTF-8"?>
<DeleteResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Deleted>
    <Key>sample1.txt</Key>
  </Deleted>
  <Error>
    <Key>sample2.txt</Key>
    <Code>AccessDenied</Code>
    <Message>Access Denied</Message>
  </Error>
</DeleteResult>
```

Sample Request for general purpose buckets: Deleting an object from a versioned bucket

If you delete an item from a versioning enabled bucket, all versions of that object remain in the bucket. However, Amazon S3 inserts a delete marker. For more information, see [Object Versioning](#).

The following scenarios describe the behavior of a multi-object Delete request when versioning is enabled for your bucket.

Case 1 - Simple Delete: In the following sample request, the multi-object delete request specifies only one key.

```
POST /?delete HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Accept: */*
x-amz-date: Wed, 30 Nov 2011 03:39:05 GMT
Content-MD5: p5/WA/oEr30qrEE121PAqw==
Authorization: AWS AKIAIOSFODNN7EXAMPLE:W0qPYCLe6JwkZAD1ei6hp9XZIIe=
Content-Length: 79
Connection: Keep-Alive

<Delete>
  <Object>
    <Key>SampleDocument.txt</Key>
  </Object>
</Delete>
```

Sample Response for general purpose buckets

Because versioning is enabled on the bucket, Amazon S3 does not delete the object. Instead, it adds a delete marker for this object. The following response indicates that a delete marker was added (the `DeleteMarker` element in the response as a value of `true`) and the version number of the delete marker it added.

```
HTTP/1.1 200 OK
x-amz-id-2: P3xqrhuhYxlrefdw3rEzmJh8z5KDtGzb+/FB7oiQaScI9Yaxd8olYXc7d1111ab+
x-amz-request-id: 264A17BF16E9E80A
Date: Wed, 30 Nov 2011 03:39:32 GMT
Content-Type: application/xml
Server: AmazonS3
Content-Length: 276

<?xml version="1.0" encoding="UTF-8"?>
<DeleteResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Deleted>
    <Key>SampleDocument.txt</Key>
    <DeleteMarker>true</DeleteMarker>
    <DeleteMarkerVersionId>NeQt5xeFTfgPJD8B4CGWnkSLtluMr11s</DeleteMarkerVersionId>
  </Deleted>
</DeleteResult>
```

Case 2 for general purpose buckets - Versioned Delete

The following request attempts to delete a specific version of an object.

```
POST /?delete HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Accept: */*
x-amz-date: Wed, 30 Nov 2011 03:39:05 GMT
Content-MD5: p5/WA/oEr30qrEE121PAqw==
Authorization: AWS AKIAIOSFODNN7EXAMPLE:W0qPYCLe6JwkZAD1ei6hp9XZIxx=
Content-Length: 140
Connection: Keep-Alive

<Delete>
  <Object>
    <Key>SampleDocument.txt</Key>
    <VersionId>0YcLXagmS.WaD..oyH4KRguB95_YhLs7</VersionId>
  </Object>
</Delete>
```

Sample Response for general purpose buckets

In this case, Amazon S3 deletes the specific object version from the bucket and returns the following response. In the response, Amazon S3 returns the key and version ID of the object deleted.

```
HTTP/1.1 400 Bad Request
x-amz-id-2: P3xqrhuYxlrefdw3rEzmJh8z5KDtGzb+/FB7oiQaScI9Yaxd8olYXc7d1111xx+
x-amz-request-id: 264A17BF16E9E80A
Date: Wed, 30 Nov 2011 03:39:32 GMT
Content-Type: application/xml
Server: AmazonS3
Content-Length: 219

<?xml version="1.0" encoding="UTF-8"?>
<DeleteResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Deleted>
    <Key>SampleDocument.txt</Key>
    <VersionId>0YcLXagmS.WaD..oyH4KRguB95_YhLs7</VersionId>
  </Deleted>
</DeleteResult>
```

Case 3 for general purpose buckets - Versioned delete of a delete marker

In the preceding example, the request refers to a delete marker (instead of an object), then Amazon S3 deletes the delete marker. The effect of this action is to make your object reappear in your bucket. Amazon S3 returns a response that indicates the delete marker it deleted (`DeleteMarker` element with value `true`) and the version ID of the delete marker.

```
HTTP/1.1 200 OK
x-amz-id-2: IIPUZrtolxDEmWsK0ae9JlSZ6yWfTye3HQ3T2iAe0ZE4XHa6NKvAJcPp51zZaBr
x-amz-request-id: D6B284CEC9B05E4E
Date: Wed, 30 Nov 2011 03:43:25 GMT
Content-Type: application/xml
Server: AmazonS3
Content-Length: 331

<?xml version="1.0" encoding="UTF-8"?>
<DeleteResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Deleted>
    <Key>SampleDocument.txt</Key>
    <VersionId>NeQt5xeFTfgPJD8B4CGWnkSLtluMr11s</VersionId>
    <DeleteMarker>true</DeleteMarker>
    <DeleteMarkerVersionId>NeQt5xeFTfgPJD8B4CGWnkSLtluMr11s</DeleteMarkerVersionId>
  </Deleted>
</DeleteResult>
```

Sample Response for general purpose buckets

In general, when a multi-object Delete request results in Amazon S3 either adding a delete marker or removing a delete marker, the response returns the following elements.

```
<DeleteMarker>true</DeleteMarker>
<DeleteMarkerVersionId>NeQt5xeFTfgPJD8B4CGWnkSLtluMr11s</DeleteMarkerVersionId>
```

Sample Request for general purpose buckets: Malformed XML in the request

This example shows how Amazon S3 responds to a request that includes a malformed XML document. The following request sends a malformed XML document (missing the Delete end element).

```
POST /?delete HTTP/1.1
```

```
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Accept: */*
x-amz-date: Wed, 30 Nov 2011 03:39:05 GMT
Content-MD5: p5/WA/oEr30qrEE121PAqw==
Authorization: AWS AKIAIOSFODNN7EXAMPLE:W0qPYCLe6JwkZAD1ei6hp9XZiIee=
Content-Length: 104
Connection: Keep-Alive
```

```
<Delete>
  <Object>
    <Key>404.txt</Key>
  </Object>
  <Object>
    <Key>a.txt</Key>
  </Object>
```

Sample Response for general purpose buckets

The response returns the error messages that describe the error.

```
HTTP/1.1 200 OK
x-amz-id-2: P3xqrihuhYxlrefdw3rEzmJh8z5KDtGzb+/FB7oiQaScI9Yaxd8olYXc7d1111ab+
x-amz-request-id: 264A17BF16E9E80A
Date: Wed, 30 Nov 2011 03:39:32 GMT
Content-Type: application/xml
Server: AmazonS3
Content-Length: 207

<?xml version="1.0" encoding="UTF-8"?>
<Error>
  <Code>MalformedXML</Code>
  <Message>The XML you provided was not well-formed or did not
    validate against our published schema</Message>
  <RequestId>264A17BF16E9E80A</RequestId>
  <HostId>P3xqrihuhYxlrefdw3rEzmJh8z5KDtGzb+/FB7oiQaScI9Yaxd8olYXc7d1111ab+</HostId>
</Error>
```

Sample Request for general purpose buckets: DeleteObjects containing a carriage return

The following example illustrates the use of an XML entity code as a substitution for a carriage return. This DeleteObjects request deletes an object with the key parameter: /some/prefix/objectwith\rcarriagereturn (where the \r is the carriage return).


```
<Delete xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
<Object>
<Key>/some/prefix/objectwith&#13;carriagereturn</Key>
</Object>
</Delete>
```

Sample Request for directory buckets: Conditional deletes

The following example shows a multi-object delete request. Each object sub-section in the XML also includes one or more preconditions around that object's Etag, Last-Modified-Time and Size that need to hold true for its deletion to succeed.

```
POST /?delete HTTP/1.1
Host: amzn-s3-demo-bucket.s3express-az_id.Region.amazonaws.com
Accept: */*
x-amz-date: Wed, 30 Nov 2011 03:39:05 GMT
Content-MD5: p5/WA/oEr30qrEE121PAqw==
Authorization: AWS AKIAIOSFODNN7EXAMPLE:W0qPYCLe6JwkZAD1ei6hp9XZiIee=
Content-Length: 125
Connection: Keep-Alive

<Delete>
<Object>
<ETag>a0e05e3566754e04b1e0f18c6b1abe1d
<LastModifiedTime>Tue, 15 Oct 2024 15:04:05 GMT</LastModifiedTime>
<Size>50</Size>
<Key>keyname1</Key>
</Object>
<Object>
<Key>keyname2</Key>
<ETag>a0e05e3566754e04b1e0f18c6b1abe1d
</Object>
<Object>
<Key>keyname3</Key>
</Object>
</Delete>
```

Sample Response for directory buckets: Conditional deletes

The response returns acknowledgments of the deleted objects.

```
HTTP/1.1 204 Success (No Content)
```

```
Content-Type: application/xml
Server: AmazonS3
x-amz-request-id: 264A17BF16E9E80A
x-amz-id-2: swzHZ0uhs7DpyZEDu
transfer-encoding: chunked
date: Wed, 30 Nov 2011 03:39:11 GMT
<?xml version="1.0" encoding="UTF-8"?>
<DeleteResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
<Deleted>
<Key>keyname1</Key>
</Deleted>
<Deleted>
<Key>keyname2</Key>
</Deleted>
<Deleted>
<Key>keyname3</Key>
</Deleted>
</DeletedResult>
```

Sample Request for general purpose buckets: Conditional deletes

The following example shows a multi-object delete request to evaluate if the objects are unchanged using the ETag value.

```
POST /?delete HTTP/1.1
Host: amzn-s3-demo-bucket.s3.Region.amazonaws.com
Accept: */*
x-amz-date: Wed, 30 Nov 2011 03:39:05 GMT
Content-MD5: p5/WA/oEr30qrEE121PAqw==
Authorization: AWS AKIAIOSFODNN7EXAMPLE:W0qPYCL6JwkZAD1ei6hp9XZIee=
Content-Length: 125
Connection: Keep-Alive

<Delete>
<Object>
<Key>keyname1</Key>
<ETag>a0e05e3566754e04b1e0f18c6b1abe1d</ETag>
</Object>
<Object>
<Key>keyname2</Key>
<ETag>a0e05e3566754e04b1e0f18c4b1abe1d</ETag>
</Object>
</Delete>
```

Sample Response for general purpose buckets: Conditional deletes

The response returns acknowledgments of the deleted objects.

```
HTTP/1.1 200 OK
Content-Type: application/xml
Server: AmazonS3
x-amz-request-id: 264A17BF16E9E80A
x-amz-id-2: swzHZ0uhs7DpyZEDu
transfer-encoding: chunked
date: Wed, 30 Nov 2011 03:39:11 GMT

<?xml version="1.0" encoding="UTF-8"?>
<DeleteResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Deleted>
    <Key>keyname1</Key>
  </Deleted>
  <Deleted>
    <Key>keyname2</Key>
  </Deleted>
</DeleteResult>
```

Sample Request for general purpose buckets: Conditional deletes

The following example shows a multi-object delete request using the If-Match header with *.

```
POST /?delete HTTP/1.1
Host: amzn-s3-demo-bucket.s3.Region.amazonaws.com
Accept: */*
x-amz-date: Wed, 30 Nov 2011 03:39:05 GMT
Content-MD5: p5/WA/oEr30qrEE121PAqw==
Authorization: AWS AKIAIOSFODNN7EXAMPLE:W0qPYCL6JwkZAD1ei6hp9XZiIee=
Content-Length: 125
Connection: Keep-Alive

<Delete>
  <Object>
    <Key>keyname1</Key>
    <ETag>*</ETag>
  </Object>
  <Object>
    <Key>keyname2</Key>
    <ETag>*</ETag>
```

```
</Object>
</Delete>
```

Sample Response for general purpose buckets: Conditional deletes

The response returns acknowledgments of the deleted objects.

```
HTTP/1.1 200 OK
Content-Type: application/xml
Server: AmazonS3
x-amz-request-id: 264A17BF16E9E80A
x-amz-id-2: swzHZ0uhs7DpyZEDu
transfer-encoding: chunked
date: Wed, 30 Nov 2011 03:39:11 GMT

<?xml version="1.0" encoding="UTF-8"?>
<DeleteResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Deleted>
    <Key>keyname1</Key>
  </Deleted>
  <Deleted>
    <Key>keyname2</Key>
  </Deleted>
</DeleteResult>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)

- [AWS SDK for Ruby V3](#)

DeleteObjectTagging

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Removes the entire tag set from the specified object. For more information about managing object tags, see [Object Tagging](#).

To use this operation, you must have permission to perform the `s3:DeleteObjectTagging` action.

To delete tags of a specific object version, add the `versionId` query parameter in the request. You will need permission for the `s3:DeleteObjectVersionTagging` action.

The following operations are related to `DeleteObjectTagging`:

- [PutObjectTagging](#)
- [GetObjectTagging](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
DELETE /{Key+}?tagging&versionId=VersionId HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name containing the objects from which to remove the tags.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form *AccessPointName-AccountId.s3-accesspoint.Region.amazonaws.com*. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

S3 on Outposts - When you use this action with S3 on Outposts, you must direct requests to the S3 on Outposts hostname. The S3 on Outposts hostname takes the form *AccessPointName-AccountId.outpostID.s3-outposts.Region.amazonaws.com*. When you use this action with S3 on Outposts, the destination bucket must be the Outposts access point ARN or the access point alias. For more information about S3 on Outposts, see [What is S3 on Outposts?](#) in the *Amazon S3 User Guide*.

Required: Yes

Key

The key that identifies the object in the bucket from which to remove all tags.

Length Constraints: Minimum length of 1.

Required: Yes

versionId

The versionId of the object that the tag-set will be removed from.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204  
x-amz-version-id: VersionId
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response.

The response returns the following HTTP headers.

x-amz-version-id

The versionId of the object the tag-set was removed from.

Examples

Sample Request

The following DELETE request deletes the tag set from the specified object.

```
DELETE /exampleobject?tagging HTTP/1.1  
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com  
Date: Wed, 25 Nov 2016 12:00:00 GMT  
Authorization: signatureValue
```

Sample Response

The following successful response shows Amazon S3 returning a 204 No Content response. The tag set for the object has been removed.

```
HTTP/1.1 204 No Content  
x-amz-version-id: VersionId  
Date: Wed, 25 Nov 2016 12:00:00 GMT
```



```
Connection: close
Server: AmazonS3
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeletePublicAccessBlock

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Removes the `PublicAccessBlock` configuration for an Amazon S3 bucket. This operation removes the bucket-level configuration only. The effective public access behavior will still be governed by account-level settings (which may inherit from organization-level policies). To use this operation, you must have the `s3:PutBucketPublicAccessBlock` permission. For more information about permissions, see [Permissions Related to Bucket Subresource Operations](#) and [Managing Access Permissions to Your Amazon S3 Resources](#).

The following operations are related to `DeletePublicAccessBlock`:

- [Using Amazon S3 Block Public Access](#)
- [GetPublicAccessBlock](#)
- [PutPublicAccessBlock](#)
- [GetBucketPolicyStatus](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
DELETE /?publicAccessBlock HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The Amazon S3 bucket whose `PublicAccessBlock` configuration you want to delete.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code `403 Forbidden` (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)

- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketAbac

Service: Amazon S3

Returns the attribute-based access control (ABAC) property of the general purpose bucket. If ABAC is enabled on your bucket, you can use tags on the bucket for access control. For more information, see [Enabling ABAC in general purpose buckets](#).

Request Syntax

```
GET /?abac HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the general purpose bucket.

Required: Yes

x-amz-expected-bucket-owner

The AWS account ID of the general purpose bucket's owner.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<AbacStatus>
  <Status>string</Status>
</AbacStatus>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

AbacStatus

Root level tag for the AbacStatus parameters.

Required: Yes

Status

The ABAC status of the general purpose bucket.

Type: String

Valid Values: Enabled | Disabled

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketAccelerateConfiguration

Service: Amazon S3

Note

This operation is not supported for directory buckets.

This implementation of the GET action uses the `accelerate` subresource to return the Transfer Acceleration state of a bucket, which is either `Enabled` or `Suspended`. Amazon S3 Transfer Acceleration is a bucket-level feature that enables you to perform faster data transfers to and from Amazon S3.

To use this operation, you must have permission to perform the `s3:GetAccelerateConfiguration` action. The bucket owner has this permission by default. The bucket owner can grant this permission to others. For more information about permissions, see [Permissions Related to Bucket Subresource Operations](#) and [Managing Access Permissions to your Amazon S3 Resources](#) in the *Amazon S3 User Guide*.

You set the Transfer Acceleration state of an existing bucket to `Enabled` or `Suspended` by using the [PutBucketAccelerateConfiguration](#) operation.

A GET `accelerate` request does not return a state value for a bucket that has no transfer acceleration state. A bucket has no Transfer Acceleration state if a state has never been set on the bucket.

For more information about transfer acceleration, see [Transfer Acceleration](#) in the Amazon S3 User Guide.

The following operations are related to `GetBucketAccelerateConfiguration`:

- [PutBucketAccelerateConfiguration](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?accelerate HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
x-amz-request-payer: RequestPayer
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket for which the accelerate configuration is retrieved.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
x-amz-request-charged: RequestCharged
<?xml version="1.0" encoding="UTF-8"?>
<AccelerateConfiguration>
  <Status>string</Status>
</AccelerateConfiguration>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

The following data is returned in XML format by the service.

AccelerateConfiguration

Root level tag for the AccelerateConfiguration parameters.

Required: Yes

Status

The accelerate configuration of the bucket.

Type: String

Valid Values: Enabled | Suspended

Examples

This implementation of the GET action returns the following responses.

Example

If the transfer acceleration state is set to Enabled on a bucket, the response is as follows:

```
<AccelerateConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Status>Enabled</Status>
</AccelerateConfiguration>
```

Example

If the transfer acceleration state is set to Suspended on a bucket, the response is as follows:

```
<AccelerateConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Status>Suspended</Status>
</AccelerateConfiguration>
```

Example

If the transfer acceleration state on a bucket has never been set to Enabled or Suspended, the response is as follows:

```
<AccelerateConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/" />
```

Retrieve the transfer acceleration configuration for a bucket

The following example shows a GET `/?accelerate` request to retrieve the transfer acceleration state of the bucket named `amzn-s3-demo-bucket`.

```
GET /?accelerate HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Mon, 11 Apr 2016 12:00:00 GMT
Authorization: authorization string
Content-Type: text/plain
```

Example

The following is a sample of the response body (only) that shows bucket transfer acceleration is enabled.

```
<AccelerateConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Status>Enabled</Status>
</AccelerateConfiguration>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketAcl

Service: Amazon S3

Note

This operation is not supported for directory buckets.

This implementation of the GET action uses the `acl` subresource to return the access control list (ACL) of a bucket. To use GET to return the ACL of the bucket, you must have the `READ_ACP` access to the bucket. If `READ_ACP` permission is granted to the anonymous user, you can return the ACL of the bucket without using an authorization header.

When you use this API operation with an access point, provide the alias of the access point in place of the bucket name.

When you use this API operation with an Object Lambda access point, provide the alias of the Object Lambda access point in place of the bucket name. If the Object Lambda access point alias in a request is not valid, the error code `InvalidAccessPointAliasError` is returned. For more information about `InvalidAccessPointAliasError`, see [List of Error Codes](#).

Note

If your bucket uses the bucket owner enforced setting for S3 Object Ownership, requests to read ACLs are still supported and return the `bucket-owner-full-control` ACL with the owner being the account that created the bucket. For more information, see [Controlling object ownership and disabling ACLs](#) in the *Amazon S3 User Guide*.

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

The following operations are related to `GetBucketAcl`:

- [ListObjects](#)

Request Syntax

```
GET /?acl HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

Specifies the S3 bucket whose ACL is being requested.

When you use this API operation with an access point, provide the alias of the access point in place of the bucket name.

When you use this API operation with an Object Lambda access point, provide the alias of the Object Lambda access point in place of the bucket name. If the Object Lambda access point alias in a request is not valid, the error code `InvalidAccessPointAliasError` is returned. For more information about `InvalidAccessPointAliasError`, see [List of Error Codes](#).

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<AccessControlPolicy>
  <Owner>
    <DisplayName>string</DisplayName>
    <ID>string</ID>
```

```
</Owner>
<AccessControlList>
  <Grant>
    <Grantee>
      <DisplayName>string</DisplayName>
      <EmailAddress>string</EmailAddress>
      <ID>string</ID>
      <xsi:type>string</xsi:type>
      <URI>string</URI>
    </Grantee>
    <Permission>string</Permission>
  </Grant>
</AccessControlList>
</AccessControlPolicy>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

AccessControlPolicy

Root level tag for the AccessControlPolicy parameters.

Required: Yes

Grants

A list of grants.

Type: Array of Grant data types

Owner

Container for the bucket owner's ID.

Type: Owner data type

Examples

Sample Request

The following request returns the ACL of the specified bucket.

```
GET ?acl HTTP/1.1
Host: bucket.s3.<Region>.amazonaws.com
Date: Wed, 28 Oct 2009 22:32:00 GMT
Authorization: authorization string
```

Sample Response

```
HTTP/1.1 200 OK
x-amz-id-2: eftixk72aD6Ap51TnqcoF8eFidJG9Z/2mkiDFu8yU9AS1ed40pIszj7UDNEHGran
x-amz-request-id: 318BC8BC148832E5
Date: Wed, 28 Oct 2009 22:32:00 GMT
Last-Modified: Sun, 1 Jan 2006 12:00:00 GMT
Content-Length: 124
Content-Type: text/plain
Connection: close
Server: AmazonS3
<AccessControlPolicy>
  <Owner>
    <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeebf76c078efc7c6caea54ba06a</ID>
  </Owner>
  <AccessControlList>
    <Grant>
      <Grantee xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:type="CanonicalUser">
        <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeebf76c078efc7c6caea54ba06a</ID>
      </Grantee>
      <Permission>FULL_CONTROL</Permission>
    </Grant>
  </AccessControlList>
</AccessControlPolicy>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)

- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketAnalyticsConfiguration

Service: Amazon S3

Note

This operation is not supported for directory buckets.

This implementation of the GET action returns an analytics configuration (identified by the analytics configuration ID) from the bucket.

To use this operation, you must have permissions to perform the `s3:GetAnalyticsConfiguration` action. The bucket owner has this permission by default. The bucket owner can grant this permission to others. For more information about permissions, see [Permissions Related to Bucket Subresource Operations](#) and [Managing Access Permissions to Your Amazon S3 Resources](#) in the *Amazon S3 User Guide*.

For information about Amazon S3 analytics feature, see [Amazon S3 Analytics – Storage Class Analysis](#) in the *Amazon S3 User Guide*.

The following operations are related to `GetBucketAnalyticsConfiguration`:

- [DeleteBucketAnalyticsConfiguration](#)
- [ListBucketAnalyticsConfigurations](#)
- [PutBucketAnalyticsConfiguration](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?analytics&id=Id HTTP/1.1
Host: Bucket.s3.amazonaws.com
```

x-amz-expected-bucket-owner: *ExpectedBucketOwner*

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket from which an analytics configuration is retrieved.

Required: Yes

id

The ID that identifies the analytics configuration.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<AnalyticsConfiguration>
  <Id>string</Id>
  <Filter>
    <And>
      <Prefix>string</Prefix>
      <Tag>
        <Key>string</Key>
        <Value>string</Value>
      </Tag>
      ...
    </And>
  </Filter>
</AnalyticsConfiguration>
```

```

<Prefix>string</Prefix>
<Tag>
  <Key>string</Key>
  <Value>string</Value>
</Tag>
</Filter>
<StorageClassAnalysis>
  <DataExport>
    <Destination>
      <S3BucketDestination>
        <Bucket>string</Bucket>
        <BucketAccountId>string</BucketAccountId>
        <Format>string</Format>
        <Prefix>string</Prefix>
      </S3BucketDestination>
    </Destination>
    <OutputSchemaVersion>string</OutputSchemaVersion>
  </DataExport>
</StorageClassAnalysis>
</AnalyticsConfiguration>

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[AnalyticsConfiguration](#)

Root level tag for the AnalyticsConfiguration parameters.

Required: Yes

[Filter](#)

The filter used to describe a set of objects for analyses. A filter must have exactly one prefix, one tag, or one conjunction (AnalyticsAndOperator). If no filter is provided, all objects will be considered in any analysis.

Type: [AnalyticsFilter](#) data type

[Id](#)

The ID that identifies the analytics configuration.

Type: String

StorageClassAnalysis

Contains data related to access patterns to be collected and made available to analyze the tradeoffs between different storage classes.

Type: [StorageClassAnalysis](#) data type

Examples

Configure an Analytics Report

The following GET request for the bucket `amzn-s3-demo-bucket` returns the inventory configuration with the ID `list1`:

```
GET /?analytics&id=list1 HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Mon, 31 Oct 2016 12:00:00 GMT
Authorization: authorization string
```

Example

The following is a sample response to the preceding GET request.

```
HTTP/1.1 200 OK
x-amz-id-2: YgIPIfBiKa2bj0KMgUAdQkf3ShJT00pXUueF6QKo
x-amz-request-id: 236A8905248E5A02
Date: Mon, 31 Oct 2016 12:00:00 GMT
Server: AmazonS3
Content-Length: length

<?xml version="1.0" encoding="UTF-8"?>
<AnalyticsConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Id>list1</Id>
  <Filter>
    <And>
      <Prefix>images/</Prefix>
      <Tag>
```

```
<Key>dog</Key>
<Value>corgi</Value>
</Tag>
</And>
</Filter>
<StorageClassAnalysis>
  <DataExport>
    <OutputSchemaVersion>V_1</OutputSchemaVersion>
    <Destination>
      <S3BucketDestination>
        <Format>CSV</Format>
        <BucketAccountId>123456789012</BucketAccountId>
        <Bucket>arn:aws:s3:::destination-bucket</Bucket>
        <Prefix>destination-prefix</Prefix>
      </S3BucketDestination>
    </Destination>
  </DataExport>
</StorageClassAnalysis>
</AnalyticsConfiguration>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketCors

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Returns the Cross-Origin Resource Sharing (CORS) configuration information set for the bucket.

To use this operation, you must have permission to perform the `s3:GetBucketCors` action. By default, the bucket owner has this permission and can grant it to others.

When you use this API operation with an access point, provide the alias of the access point in place of the bucket name.

When you use this API operation with an Object Lambda access point, provide the alias of the Object Lambda access point in place of the bucket name. If the Object Lambda access point alias in a request is not valid, the error code `InvalidAccessPointAliasError` is returned. For more information about `InvalidAccessPointAliasError`, see [List of Error Codes](#).

For more information about CORS, see [Enabling Cross-Origin Resource Sharing](#).

The following operations are related to `GetBucketCors`:

- [PutBucketCors](#)
- [DeleteBucketCors](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?cors HTTP/1.1
```

```
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name for which to get the cors configuration.

When you use this API operation with an access point, provide the alias of the access point in place of the bucket name.

When you use this API operation with an Object Lambda access point, provide the alias of the Object Lambda access point in place of the bucket name. If the Object Lambda access point alias in a request is not valid, the error code `InvalidAccessPointAliasError` is returned. For more information about `InvalidAccessPointAliasError`, see [List of Error Codes](#).

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<CORSConfiguration>
  <CORSRule>
    <AllowedHeader>string</AllowedHeader>
    ...
    <AllowedMethod>string</AllowedMethod>
    ...
```

```
<AllowedOrigin>string</AllowedOrigin>
...
<ExposeHeader>string</ExposeHeader>
...
<ID>string</ID>
<MaxAgeSeconds>integer</MaxAgeSeconds>
</CORSRule>
...
</CORSConfiguration>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[CORSConfiguration](#)

Root level tag for the CORSConfiguration parameters.

Required: Yes

[CORSRule](#)

A set of origins and methods (cross-origin access that you want to allow). You can add up to 100 rules to the configuration.

Type: Array of [CORSRule](#) data types

Examples

Configure CORS Sample Request

The following PUT request adds the cors subresource to a bucket (amzn-s3-demo-bucket).

```
PUT /?cors HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
x-amz-date: Tue, 21 Aug 2012 17:54:50 GMT
Content-MD5: 8dYiLewFWZyGgV2Q5FNI4W==
Authorization: authorization string
Content-Length: 216
```



```
<CORSConfiguration>
  <CORSRule>
    <AllowedOrigin>http://www.example.com</AllowedOrigin>
    <AllowedMethod>PUT</AllowedMethod>
    <AllowedMethod>POST</AllowedMethod>
    <AllowedMethod>DELETE</AllowedMethod>
    <AllowedHeader>*</AllowedHeader>
    <MaxAgeSeconds>3000</MaxAgeSec>
    <ExposeHeader>x-amz-server-side-encryption</ExposeHeader>
  </CORSRule>
  <CORSRule>
    <AllowedOrigin>*</AllowedOrigin>
    <AllowedMethod>GET</AllowedMethod>
    <AllowedHeader>*</AllowedHeader>
    <MaxAgeSeconds>3000</MaxAgeSeconds>
  </CORSRule>
</CORSConfiguration>
```

Example

This is the sample response to the preceding request.

```
HTTP/1.1 200 OK
x-amz-id-2: CCsh0vb0Pfzxzhw0ADyC4qHj/Ck3F9Q0viXKw3rivZ+GcBoZS00ahvEJfPisZB7B
x-amz-request-id: BDC4B83DF5096BBE
Date: Tue, 21 Aug 2012 17:54:50 GMT
Server: AmazonS3
```

Sample Request: Retrieve cors subresource

The following example gets the cors subresource of a bucket.

```
GET /?cors HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Tue, 13 Dec 2011 19:14:42 GMT
Authorization: signatureValue
```

Example

Sample Response

```
HTTP/1.1 200 OK
x-amz-id-2: 0FmFIWsh/
PpBuzZ0JFRC55ZGVmQW4SHJ7xVDqKwhEdJmf3q63RtrvH8ZuxW1Bo15
x-amz-request-id: 0CF038E9BCF63097
Date: Tue, 13 Dec 2011 19:14:42 GMT
Server: AmazonS3
Content-Length: 280
<CORSConfiguration>
  <CORSRule>
    <AllowedOrigin>http://www.example.com</AllowedOrigin>
    <AllowedMethod>GET</AllowedMethod>
    <MaxAgeSeconds>3000</MaxAgeSec>
    <ExposeHeader>x-amz-server-side-encryption</ExposeHeader>
  </CORSRule>
</CORSConfiguration>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketEncryption

Service: Amazon S3

Returns the default encryption configuration for an Amazon S3 bucket. By default, all buckets have a default encryption configuration that uses server-side encryption with Amazon S3 managed keys (SSE-S3). This operation also returns the [BucketKeyEnabled](#) and [BlockedEncryptionTypes](#) statuses.

Note

- **General purpose buckets** - For information about the bucket default encryption feature, see [Amazon S3 Bucket Default Encryption](#) in the *Amazon S3 User Guide*.
- **Directory buckets** - For directory buckets, there are only two supported options for server-side encryption: SSE-S3 and SSE-KMS. For information about the default encryption configuration in directory buckets, see [Setting default server-side encryption behavior for directory buckets](#).

Permissions

- **General purpose bucket permissions** - The `s3:GetEncryptionConfiguration` permission is required in a policy. The bucket owner has this permission by default. The bucket owner can grant this permission to others. For more information about permissions, see [Permissions Related to Bucket Operations](#) and [Managing Access Permissions to Your Amazon S3 Resources](#).
- **Directory bucket permissions** - To grant access to this API operation, you must have the `s3express:GetEncryptionConfiguration` permission in an IAM identity-based policy instead of a bucket policy. Cross-account access to this API operation isn't supported. This operation can only be performed by the AWS account that owns the resource. For more information about directory bucket policies and permissions, see [AWS Identity and Access Management \(IAM\) for S3 Express One Zone](#) in the *Amazon S3 User Guide*.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `s3express-control.region-code.amazonaws.com`.

The following operations are related to `GetBucketEncryption`:

- [PutBucketEncryption](#)
- [DeleteBucketEncryption](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?encryption HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

[Bucket](#)

The name of the bucket from which the server-side encryption configuration is retrieved.

Directory buckets - When you use this operation with a directory bucket, you must use path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name` . Virtual-hosted-style requests aren't supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must also follow the format `bucket-base-name--zone-id--x-s3` (for example, `DOC-EXAMPLE-BUCKET--usw2-az1--x-s3`). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*

Required: Yes

[x-amz-expected-bucket-owner](#)

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Note

For directory buckets, this header is not supported in this API operation. If you specify this header, the request fails with the HTTP status code 501 Not Implemented.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<ServerSideEncryptionConfiguration>
  <Rule>
    <ApplyServerSideEncryptionByDefault>
      <KMSEMasterKeyID>string</KMSEMasterKeyID>
      <SSEAlgorithm>string</SSEAlgorithm>
    </ApplyServerSideEncryptionByDefault>
    <BlockedEncryptionTypes>
      <EncryptionType>string</EncryptionType>
      ...
    </BlockedEncryptionTypes>
    <BucketKeyEnabled>boolean</BucketKeyEnabled>
  </Rule>
  ...
</ServerSideEncryptionConfiguration>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

ServerSideEncryptionConfiguration

Root level tag for the ServerSideEncryptionConfiguration parameters.

Required: Yes

Rule

Container for information about a particular server-side encryption configuration rule.

Type: Array of [ServerSideEncryptionRule](#) data types

Examples

Sample Request: Retrieve the encryption configuration for an S3 general purpose bucket

The following example shows a GET `/?encryption` request.

```
GET /?encryption HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Wed, 06 Sep 2017 12:00:00 GMT
Authorization: authorization string
Content-Length: length
```

Sample Response for a general purpose bucket

This example illustrates one usage of `GetBucketEncryption`.

```
HTTP/1.1 200 OK
x-amz-id-2: kDmqsuw5FDmgLmxQaUkd9A4NJ/PIiE0c1rAU/ue2Yp60toXs4I5k5fqlwZsA6fV
+wJQCzRRwygQ=
x-amz-request-id: 5D8706FCB2673B7D
Date: Wed, 19 Nov 2025 12:00:00 GMT
Transfer-Encoding: chunked
Server: AmazonS3

<ServerSideEncryptionConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Rule>
    <ApplyServerSideEncryptionByDefault>
      <SSEAlgorithm>aws:kms</SSEAlgorithm>
      <KMSKeyID>arn:aws:kms:us-east-1:1234/5678example</KMSKeyID>
    </ApplyServerSideEncryptionByDefault>
    <BlockedEncryptionTypes>
      <EncryptionType>SSE-C</EncryptionType>
    </BlockedEncryptionTypes>
  </Rule>
</ServerSideEncryptionConfiguration>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketIntelligentTieringConfiguration

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Gets the S3 Intelligent-Tiering configuration from the specified bucket.

The S3 Intelligent-Tiering storage class is designed to optimize storage costs by automatically moving data to the most cost-effective storage access tier, without performance impact or operational overhead. S3 Intelligent-Tiering delivers automatic cost savings in three low latency and high throughput access tiers. To get the lowest storage cost on data that can be accessed in minutes to hours, you can choose to activate additional archiving capabilities.

The S3 Intelligent-Tiering storage class is the ideal storage class for data with unknown, changing, or unpredictable access patterns, independent of object size or retention period. If the size of an object is less than 128 KB, it is not monitored and not eligible for auto-tiering. Smaller objects can be stored, but they are always charged at the Frequent Access tier rates in the S3 Intelligent-Tiering storage class.

For more information, see [Storage class for automatically optimizing frequently and infrequently accessed objects](#).

Operations related to GetBucketIntelligentTieringConfiguration include:

- [DeleteBucketIntelligentTieringConfiguration](#)
- [PutBucketIntelligentTieringConfiguration](#)
- [ListBucketIntelligentTieringConfigurations](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?intelligent-tiering&id=Id HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the Amazon S3 bucket whose configuration you want to modify or retrieve.

Required: Yes

id

The ID used to identify the S3 Intelligent-Tiering configuration.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<IntelligentTieringConfiguration>
  <Id>string</Id>
  <Filter>
    <And>
      <Prefix>string</Prefix>
      <Tag>
        <Key>string</Key>
```

```

        <Value>string</Value>
      </Tag>
      ...
    </And>
    <Prefix>string</Prefix>
    <Tag>
      <Key>string</Key>
      <Value>string</Value>
    </Tag>
  </Filter>
  <Status>string</Status>
  <Tiering>
    <AccessTier>string</AccessTier>
    <Days>integer</Days>
  </Tiering>
  ...
</IntelligentTieringConfiguration>

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

IntelligentTieringConfiguration

Root level tag for the IntelligentTieringConfiguration parameters.

Required: Yes

Filter

Specifies a bucket filter. The configuration only includes objects that meet the filter's criteria.

Type: [IntelligentTieringFilter](#) data type

Id

The ID used to identify the S3 Intelligent-Tiering configuration.

Type: String

Status

Specifies the status of the configuration.

Type: String

Valid Values: Enabled | Disabled

Tiering

Specifies the S3 Intelligent-Tiering storage class tier of the configuration.

Type: Array of [Tiering](#) data types

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketInventoryConfiguration

Service: Amazon S3

Returns an S3 Inventory configuration (identified by the inventory configuration ID) from the bucket.

Note

Directory buckets - For directory buckets, you must make requests for this API operation to the Regional endpoint. These endpoints support path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name` . Virtual-hosted-style requests aren't supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

Permissions

To use this operation, you must have permissions to perform the `s3:GetInventoryConfiguration` action. The bucket owner has this permission by default. The bucket owner can grant this permission to others. For more information about permissions, see [Permissions Related to Bucket Subresource Operations](#) and [Managing Access Permissions to Your Amazon S3 Resources](#).

- **General purpose bucket permissions** - The `s3:GetInventoryConfiguration` permission is required in a policy. For more information about general purpose buckets permissions, see [Using Bucket Policies and User Policies](#) in the *Amazon S3 User Guide*.
- **Directory bucket permissions** - To grant access to this API operation, you must have the `s3express:GetInventoryConfiguration` permission in an IAM identity-based policy instead of a bucket policy. For more information about directory bucket policies and permissions, see [AWS Identity and Access Management \(IAM\) for S3 Express One Zone](#) in the *Amazon S3 User Guide*.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `s3express-control.region-code.amazonaws.com`.

For information about the Amazon S3 inventory feature, see [Amazon S3 Inventory](#).

The following operations are related to `GetBucketInventoryConfiguration`:

- [DeleteBucketInventoryConfiguration](#)
- [ListBucketInventoryConfigurations](#)
- [PutBucketInventoryConfiguration](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?inventory&id=Id HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket containing the inventory configuration to retrieve.

Directory buckets - When you use this operation with a directory bucket, you must use path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name`. Virtual-hosted-style requests aren't supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must also follow the format `bucket-base-name--zone-id--x-s3` (for example, `DOC-EXAMPLE-BUCKET--usw2-az1--x-s3`). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*.

Required: Yes

id

The ID used to identify the inventory configuration.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Note

For directory buckets, this header is not supported in this API operation. If you specify this header, the request fails with the HTTP status code 501 Not Implemented.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<InventoryConfiguration>
  <Destination>
    <S3BucketDestination>
      <AccountId>string</AccountId>
      <Bucket>string</Bucket>
      <Encryption>
        <SSE-KMS>
          <KeyId>string</KeyId>
        </SSE-KMS>
        <SSE-S3>
          </SSE-S3>
        </Encryption>
      <Format>string</Format>
      <Prefix>string</Prefix>
    </S3BucketDestination>
  </Destination>
```

```
<IsEnabled>boolean</IsEnabled>
<Filter>
  <Prefix>string</Prefix>
</Filter>
<Id>string</Id>
<IncludedObjectVersions>string</IncludedObjectVersions>
<OptionalFields>
  <Field>string</Field>
</OptionalFields>
<Schedule>
  <Frequency>string</Frequency>
</Schedule>
</InventoryConfiguration>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

InventoryConfiguration

Root level tag for the InventoryConfiguration parameters.

Required: Yes

Destination

Contains information about where to publish the inventory results.

Type: [InventoryDestination](#) data type

Filter

Specifies an inventory filter. The inventory only includes objects that meet the filter's criteria.

Type: [InventoryFilter](#) data type

Id

The ID used to identify the inventory configuration.

Type: String

IncludedObjectVersions

Object versions to include in the inventory list. If set to `All`, the list includes all the object versions, which adds the version-related fields `VersionId`, `IsLatest`, and `DeleteMarker` to the list. If set to `Current`, the list does not contain these version-related fields.

Type: String

Valid Values: `All` | `Current`

IsEnabled

Specifies whether the inventory is enabled or disabled. If set to `True`, an inventory list is generated. If set to `False`, no inventory list is generated.

Type: Boolean

OptionalFields

Contains the optional fields that are included in the inventory results.

Note

The following optional fields are supported for directory buckets `Size` | `LastModifiedDate` | `StorageClass` | `ETag` | `IsMultipartUploaded` | `EncryptionStatus` | `BucketKeyStatus` | `ChecksumAlgorithm` | `LifecycleExpirationDate`. Throws `MalformedXML` error if unsupported optional field is provided.

Type: Array of strings

Valid Values: `Size` | `LastModifiedDate` | `StorageClass` | `ETag` | `IsMultipartUploaded` | `ReplicationStatus` | `EncryptionStatus` | `ObjectLockRetainUntilDate` | `ObjectLockMode` | `ObjectLockLegalHoldStatus` | `ObjectLockEventHoldStatus` | `ObjectLockEventHoldDuration` | `IntelligentTieringAccessTier` | `BucketKeyStatus` | `ChecksumAlgorithm` | `ObjectAccessControlList` | `ObjectOwner` | `LifecycleExpirationDate`

Schedule

Specifies the schedule for generating inventory results.

Type: [InventorySchedule](#) data type

Examples

Sample Request: Configure an inventory report

The following GET request for the bucket `examplebucket` returns the inventory configuration with the ID `list1`.

```
GET /?inventory&id=list1 HTTP/1.1
Host: examplebucket.s3.<Region>.amazonaws.com
Date: Mon, 31 Oct 2016 12:00:00 GMT
Authorization: authorization string
```

Sample Response

This example illustrates one usage of `GetBucketInventoryConfiguration`.

```
HTTP/1.1 200 OK
x-amz-id-2: YgIPIfBiKa2bj0KMgUAdQkf3ShJT00pXUueF6QKo
x-amz-request-id: 236A8905248E5A02
Date: Mon, 31 Oct 2016 12:00:00 GMT
Server: AmazonS3
Content-Length: length

<?xml version="1.0" encoding="UTF-8"?>
<InventoryConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Id>report1</Id>
  <IsEnabled>true</IsEnabled>
  <Destination>
    <S3BucketDestination>
      <Format>CSV</Format>
      <AccountId>123456789012</AccountId>
      <Bucket>arn:aws:s3:::destination-bucket</Bucket>
      <Prefix>prefix1</Prefix>
      <SSE-S3/>
    </S3BucketDestination>
  </Destination>
  <Schedule>
```

```
        <Frequency>Daily</Frequency>
    </Schedule>
    <Filter>
        <Prefix>myprefix</Prefix>
    </Filter>
    <IncludedObjectVersions>All</IncludedObjectVersions>
    <OptionalFields>
        <Field>Size</Field>
        <Field>LastModifiedDate</Field>
        <Field>ETag</Field>
        <Field>StorageClass</Field>
        <Field>IsMultipartUploaded</Field>
        <Field>ReplicationStatus</Field>
        <Field>ObjectLockRetainUntilDate</Field>
        <Field>ObjectLockMode</Field>
        <Field>ObjectLockLegalHoldStatus</Field>
    </OptionalFields>
</InventoryConfiguration>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketLifecycle

Service: Amazon S3

Important

For an updated version of this API, see [GetBucketLifecycleConfiguration](#). If you configured a bucket lifecycle using the `filter` element, you should see the updated version of this topic. This topic is provided for backward compatibility.

Note

This operation is not supported for directory buckets.

Returns the lifecycle configuration information set on the bucket. For information about lifecycle configuration, see [Object Lifecycle Management](#).

To use this operation, you must have permission to perform the `s3:GetLifecycleConfiguration` action. The bucket owner has this permission by default. The bucket owner can grant this permission to others. For more information about permissions, see [Permissions Related to Bucket Subresource Operations](#) and [Managing Access Permissions to Your Amazon S3 Resources](#).

GetBucketLifecycle has the following special error:

- Error code: `NoSuchLifecycleConfiguration`
 - Description: The lifecycle configuration does not exist.
 - HTTP Status Code: 404 Not Found
 - SOAP Fault Code Prefix: Client

The following operations are related to GetBucketLifecycle:

- [GetBucketLifecycleConfiguration](#)
- [PutBucketLifecycle](#)
- [DeleteBucketLifecycle](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?lifecycle HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket for which to get the lifecycle information.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<LifecycleConfiguration>
  <Rule>
    <AbortIncompleteMultipartUpload>
      <DaysAfterInitiation>integer</DaysAfterInitiation>
```

```

</AbortIncompleteMultipartUpload>
<Expiration>
  <Date>timestamp</Date>
  <Days>integer</Days>
  <ExpiredObjectDeleteMarker>boolean</ExpiredObjectDeleteMarker>
</Expiration>
<ID>string</ID>
<NoncurrentVersionExpiration>
  <NewerNoncurrentVersions>integer</NewerNoncurrentVersions>
  <NoncurrentDays>integer</NoncurrentDays>
</NoncurrentVersionExpiration>
<NoncurrentVersionTransition>
  <NewerNoncurrentVersions>integer</NewerNoncurrentVersions>
  <NoncurrentDays>integer</NoncurrentDays>
  <StorageClass>string</StorageClass>
</NoncurrentVersionTransition>
<Prefix>string</Prefix>
<Status>string</Status>
<Transition>
  <Date>timestamp</Date>
  <Days>integer</Days>
  <StorageClass>string</StorageClass>
</Transition>
</Rule>
...
</LifecycleConfiguration>

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

LifecycleConfiguration

Root level tag for the LifecycleConfiguration parameters.

Required: Yes

Rule

Container for a lifecycle rule.

Type: Array of [Rule](#) data types

Examples

Sample Request: Retrieve a lifecycle subresource

This example is a GET request to retrieve the lifecycle subresource from the specified bucket, and an example response with the returned lifecycle configuration.

```
GET /?lifecycle HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
x-amz-date: Thu, 15 Nov 2012 00:17:21 GMT
Authorization: signatureValue
```

Sample Response

This example illustrates one usage of GetBucketLifecycle.

```
HTTP/1.1 200 OK
x-amz-id-2:
ITnGT1y4RyTmXa3rPi4hklTXouTf0hccUjo0iCPjz6FnfIutBj3M7fPGlW02SEWp
x-amz-request-id: 51991C342C575321
Date: Thu, 15 Nov 2012 00:17:23 GMT
Server: AmazonS3
Content-Length: 358

<?xml version="1.0" encoding="UTF-8"?>
<LifecycleConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Rule>
    <ID>Archive and then delete rule</ID>
    <Prefix>projectdocs</Prefix>
    <Status>Enabled</Status>
    <Transition>
      <Days>30</Days>
      <StorageClass>STANDARD_IA</StorageClass>
    </Transition>
    <Transition>
      <Days>365</Days>
      <StorageClass>GLACIER</StorageClass>
    </Transition>
    <Expiration>
      <Days>3650</Days>
```

```
        </Expiration>
    </Rule>
</LifecycleConfiguration>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketLifecycleConfiguration

Service: Amazon S3

Returns the lifecycle configuration information set on the bucket. For information about lifecycle configuration, see [Object Lifecycle Management](#).

Bucket lifecycle configuration now supports specifying a lifecycle rule using an object key name prefix, one or more object tags, object size, or any combination of these. Accordingly, this section describes the latest API, which is compatible with the new functionality. The previous version of the API supported filtering based only on an object key name prefix, which is supported for general purpose buckets for backward compatibility. For the related API description, see [GetBucketLifecycle](#).

Note

Lifecycle configurations for directory buckets only support expiring objects and cancelling multipart uploads. Expiring of versioned objects, transitions and tag filters are not supported.

Permissions

- **General purpose bucket permissions** - By default, all Amazon S3 resources are private, including buckets, objects, and related subresources (for example, lifecycle configuration and website configuration). Only the resource owner (that is, the AWS account that created it) can access the resource. The resource owner can optionally grant access permissions to others by writing an access policy. For this operation, a user must have the `s3:GetLifecycleConfiguration` permission.

For more information about permissions, see [Managing Access Permissions to Your Amazon S3 Resources](#).

- **Directory bucket permissions** - You must have the `s3express:GetLifecycleConfiguration` permission in an IAM identity-based policy to use this operation. Cross-account access to this API operation isn't supported. The resource owner can optionally grant access permissions to others by creating a role or user for them as long as they are within the same account as the owner and resource.

For more information about directory bucket policies and permissions, see [Authorizing Regional endpoint APIs with IAM](#) in the *Amazon S3 User Guide*.

 Note

Directory buckets - For directory buckets, you must make requests for this API operation to the Regional endpoint. These endpoints support path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name`. Virtual-hosted-style requests aren't supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `s3express-control.region.amazonaws.com`.

`GetBucketLifecycleConfiguration` has the following special error:

- Error code: `NoSuchLifecycleConfiguration`
- Description: The lifecycle configuration does not exist.
- HTTP Status Code: 404 Not Found
- SOAP Fault Code Prefix: Client

The following operations are related to `GetBucketLifecycleConfiguration`:

- [GetBucketLifecycle](#)
- [PutBucketLifecycle](#)
- [DeleteBucketLifecycle](#)

 Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?lifecycle HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket for which to get the lifecycle information.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Note

This parameter applies to general purpose buckets only. It is not supported for directory bucket lifecycle configurations.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
x-amz-transition-default-minimum-object-size: TransitionDefaultMinimumObjectSize
<?xml version="1.0" encoding="UTF-8"?>
<LifecycleConfiguration>
  <Rule>
    <AbortIncompleteMultipartUpload>
      <DaysAfterInitiation>integer</DaysAfterInitiation>
    </AbortIncompleteMultipartUpload>
```

```

<Expiration>
  <Date>timestamp</Date>
  <Days>integer</Days>
  <ExpiredObjectDeleteMarker>boolean</ExpiredObjectDeleteMarker>
</Expiration>
<Filter>
  <And>
    <ObjectSizeGreaterThan>long</ObjectSizeGreaterThan>
    <ObjectSizeLessThan>long</ObjectSizeLessThan>
    <Prefix>string</Prefix>
    <Tag>
      <Key>string</Key>
      <Value>string</Value>
    </Tag>
    ...
  </And>
  <ObjectSizeGreaterThan>long</ObjectSizeGreaterThan>
  <ObjectSizeLessThan>long</ObjectSizeLessThan>
  <Prefix>string</Prefix>
  <Tag>
    <Key>string</Key>
    <Value>string</Value>
  </Tag>
</Filter>
<ID>string</ID>
<NoncurrentVersionExpiration>
  <NewerNoncurrentVersions>integer</NewerNoncurrentVersions>
  <NoncurrentDays>integer</NoncurrentDays>
</NoncurrentVersionExpiration>
<NoncurrentVersionTransition>
  <NewerNoncurrentVersions>integer</NewerNoncurrentVersions>
  <NoncurrentDays>integer</NoncurrentDays>
  <StorageClass>string</StorageClass>
</NoncurrentVersionTransition>
...
<Prefix>string</Prefix>
<Status>string</Status>
<Transition>
  <Date>timestamp</Date>
  <Days>integer</Days>
  <StorageClass>string</StorageClass>
</Transition>
...
</Rule>

```

```
...  
</LifecycleConfiguration>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

x-amz-transition-default-minimum-object-size

Indicates which default minimum object size behavior is applied to the lifecycle configuration.

Note

This parameter applies to general purpose buckets only. It isn't supported for directory bucket lifecycle configurations.

- `all_storage_classes_128K` - Objects smaller than 128 KB will not transition to any storage class by default.
- `varies_by_storage_class` - Objects smaller than 128 KB will transition to Glacier Flexible Retrieval or Glacier Deep Archive storage classes. By default, all other storage classes will prevent transitions smaller than 128 KB.

To customize the minimum object size for any transition you can add a filter that specifies a custom `ObjectSizeGreaterThan` or `ObjectSizeLessThan` in the body of your transition rule. Custom filters always take precedence over the default transition behavior.

Valid Values: `varies_by_storage_class` | `all_storage_classes_128K`

The following data is returned in XML format by the service.

LifecycleConfiguration

Root level tag for the LifecycleConfiguration parameters.

Required: Yes

Rule

Container for a lifecycle rule.

Type: Array of [LifecycleRule](#) data types

Examples

Example 1: Get lifecycle configuration - general purpose bucket

This example illustrates how to use `GetBucketLifecycleConfiguration` to retrieve the lifecycle configuration for a general purpose bucket:

```
GET /?lifecycle HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
x-amz-date: Thu, 15 Nov 2012 00:17:21 GMT
Authorization: signatureValue
```

Sample Response

This example shows the response from the preceding `GetBucketLifecycleConfiguration` request:

```
HTTP/1.1 200 OK
x-amz-id-2: ITnGT1y4RyTmXa3rPi4hklTXouTf0hccUjo0iCPjz6FnfIutBj3M7fPGlW02SEWp
x-amz-request-id: 51991C342C575321
Date: Thu, 15 Nov 2012 00:17:23 GMT
Server: AmazonS3
Content-Length: 358

<?xml version="1.0" encoding="UTF-8"?>
<LifecycleConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Rule>
    <ID>Archive and then delete rule</ID>
    <Prefix>projectdocs</Prefix>
    <Status>Enabled</Status>
    <Transition>
      <Days>30</Days>
      <StorageClass>STANDARD_IA</StorageClass>
    </Transition>
    <Transition>
      <Days>365</Days>
      <StorageClass>GLACIER</StorageClass>
    </Transition>
    <Expiration>
```

```
        <Days>3650</Days>
    </Expiration>
</Rule>
</LifecycleConfiguration>
```

Example 2: Get lifecycle configuration - directory bucket

This example illustrates how to use `GetBucketLifecycleConfiguration` to retrieve the lifecycle configuration for a directory bucket:

```
GET /?lifecycle HTTP/1.1
Host:s3express-control.us-west-2.amazonaws.com
```

Sample Response

This example shows the response from the preceding `GetBucketLifecycleConfiguration` request:

```
<?xml version="1.0" encoding="UTF-8"?>
<LifecycleConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Rule>
    <ID>Lifecycle expiration rule</ID>
    <Filter>
      <And>
        <Prefix>myprefix</Prefix>
        <ObjectSizeGreaterThan>500</ObjectSizeGreaterThan>
        <ObjectSizeLessThan>64000</ObjectSizeLessThan>
      </And>
    </Filter>
    <Status>Enabled</Status>
    <Expiration>
      <Days>7</Days>
    </Expiration>
  </Rule>
  <Rule>
    <ID>MPU Rule </ID>
    <Filter>
      <Prefix>another_prefix </Prefix>
    </Filter>
```

```
<Status>Enabled</Status>
<AbortIncompleteMultipartUpload>
  <DaysAfterInitiation>3</DaysAfterInitiation>
</AbortIncompleteMultipartUpload>
</Rule>
</LifecycleConfiguration>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketLocation

Service: Amazon S3

Important

Using the `GetBucketLocation` operation is no longer a best practice. To return the Region that a bucket resides in, we recommend that you use the [HeadBucket](#) operation instead. For backward compatibility, Amazon S3 continues to support the `GetBucketLocation` operation.

Returns the Region the bucket resides in. You set the bucket's Region using the `LocationConstraint` request parameter in a `CreateBucket` request. For more information, see [CreateBucket](#).

Note

In a bucket's home Region, calls to the `GetBucketLocation` operation are governed by the bucket's policy. In other Regions, the bucket policy doesn't apply, which means that cross-account access won't be authorized. However, calls to the `HeadBucket` operation always return the bucket's location through an HTTP response header, whether access to the bucket is authorized or not. Therefore, we recommend using the `HeadBucket` operation for bucket Region discovery and to avoid using the `GetBucketLocation` operation.

When you use this API operation with an access point, provide the alias of the access point in place of the bucket name.

When you use this API operation with an Object Lambda access point, provide the alias of the Object Lambda access point in place of the bucket name. If the Object Lambda access point alias in a request is not valid, the error code `InvalidAccessPointAliasError` is returned. For more information about `InvalidAccessPointAliasError`, see [List of Error Codes](#).

Note

This operation is not supported for directory buckets.

The following operations are related to `GetBucketLocation`:

- [GetObject](#)
- [CreateBucket](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?location HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket for which to get the location.

When you use this API operation with an access point, provide the alias of the access point in place of the bucket name.

When you use this API operation with an Object Lambda access point, provide the alias of the Object Lambda access point in place of the bucket name. If the Object Lambda access point alias in a request is not valid, the error code `InvalidAccessPointAliasError` is returned. For more information about `InvalidAccessPointAliasError`, see [List of Error Codes](#).

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code `403 Forbidden` (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<LocationConstraint>
  <LocationConstraint>string</LocationConstraint>
</LocationConstraint>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

LocationConstraint

Root level tag for the LocationConstraint parameters.

Required: Yes

LocationConstraint

Specifies the Region where the bucket resides. For a list of all the Amazon S3 supported location constraints by Region, see [Regions and Endpoints](#).

Buckets in Region us-east-1 have a LocationConstraint of null. Buckets with a LocationConstraint of EU reside in eu-west-1.

Type: String

Valid Values: af-south-1 | ap-east-1 | ap-east-2 | ap-northeast-1 | ap-northeast-2 | ap-northeast-3 | ap-south-1 | ap-south-2 | ap-southeast-1 | ap-southeast-2 | ap-southeast-3 | ap-southeast-4 | ap-southeast-5 | ap-southeast-6 | ap-southeast-7 | ca-central-1 | ca-west-1 | cn-north-1 | cn-northwest-1 | EU | eu-central-1 | eu-central-2 | eu-north-1 | eu-south-1 | eu-south-2 | eu-west-1 | eu-west-2 | eu-west-3 | il-central-1 | me-central-1 | me-south-1 | mx-central-1 | sa-east-1 | us-east-2 | us-gov-east-1 | us-gov-west-1 | us-west-1 | us-west-2

Examples

Sample Request

The following request returns the Region of the specified bucket.

```
GET /?location HTTP/1.1
Host: amzn-s3-demo-bucket.s3.amazonaws.com
Date: Tue, 09 Oct 2007 20:26:04 +0000
Authorization: signatureValue
```

Sample Response

This example illustrates one usage of GetBucketLocation.

```
<?xml version="1.0" encoding="UTF-8"?>
  <LocationConstraint xmlns="http://s3.amazonaws.com/doc/2006-03-01/">us-
west-2</LocationConstraint>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketLogging

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Returns the logging status of a bucket and the permissions users have to view and modify that status.

The following operations are related to GetBucketLogging:

- [CreateBucket](#)
- [PutBucketLogging](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?logging HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name for which to get the logging information.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<BucketLoggingStatus>
  <LoggingEnabled>
    <TargetBucket>string</TargetBucket>
    <TargetGrants>
      <Grant>
        <Grantee>
          <DisplayName>string</DisplayName>
          <EmailAddress>string</EmailAddress>
          <ID>string</ID>
          <xsi:type>string</xsi:type>
          <URI>string</URI>
        </Grantee>
        <Permission>string</Permission>
      </Grant>
    </TargetGrants>
    <TargetObjectKeyFormat>
      <PartitionedPrefix>
        <PartitionDateSource>string</PartitionDateSource>
      </PartitionedPrefix>
      <SimplePrefix>
      </SimplePrefix>
    </TargetObjectKeyFormat>
    <TargetPrefix>string</TargetPrefix>
  </LoggingEnabled>
</BucketLoggingStatus>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

BucketLoggingStatus

Root level tag for the BucketLoggingStatus parameters.

Required: Yes

LoggingEnabled

Describes where logs are stored and the prefix that Amazon S3 assigns to all log object keys for a bucket. For more information, see [PUT Bucket logging](#) in the *Amazon S3 API Reference*.

Type: [LoggingEnabled](#) data type

Examples

Sample Request

The following request returns the logging status for amzn-s3-demo-bucket.

```
GET ?logging HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Wed, 25 Nov 2009 12:00:00 GMT
Authorization: authorization string
```

Sample Response: Showing an enabled logging status

This example illustrates one usage of GetBucketLogging.

```
HTTP/1.1 200 OK
Date: Wed, 25 Nov 2009 12:00:00 GMT
Connection: close
Server: AmazonS3

<?xml version="1.0" encoding="UTF-8"?>
```

```

<BucketLoggingStatus xmlns="http://doc.s3.amazonaws.com/2006-03-01">
  <LoggingEnabled>
    <TargetBucket>amzn-s3-demo-bucket</TargetBucket>
    <TargetPrefix>mybucket-access_log-</TargetPrefix>
    <TargetGrants>
      <Grant>
        <Grantee xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
          xsi:type="CanonicalUser">
          <ID>79a59df900b949e55d96a1e698fbacedfd6e09d98eacf8f8d5218e7cd47ef2be</ID>
          </Grantee>
          <Permission>READ</Permission>
        </Grant>
      </TargetGrants>
    </LoggingEnabled>
  </BucketLoggingStatus>

```

Sample Response: Showing a disabled logging status

This example illustrates one usage of GetBucketLogging.

```

HTTP/1.1 200 OK
Date: Wed, 25 Nov 2009 12:00:00 GMT
Connection: close
Server: AmazonS3

<?xml version="1.0" encoding="UTF-8"?>

<BucketLoggingStatus xmlns="http://doc.s3.amazonaws.com/2006-03-01">
  <!--<LoggingEnabled><TargetBucket>myLogsBucket</
TargetBucket><TargetPrefix>add/this/prefix/to/my/log/files/access_log-</TargetPrefix></
LoggingEnabled>-->
</BucketLoggingStatus>

```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)

- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketMetadataConfiguration

Service: Amazon S3

Retrieves the S3 Metadata configuration for a general purpose bucket. For more information, see [Accelerating data discovery with S3 Metadata](#) in the *Amazon S3 User Guide*.

Note

You can use the V2 GetBucketMetadataConfiguration API operation with V1 or V2 metadata configurations. However, if you try to use the V1 GetBucketMetadataTableConfiguration API operation with V2 configurations, you will receive an HTTP 405 Method Not Allowed error.

Permissions

To use this operation, you must have the `s3:GetBucketMetadataTableConfiguration` permission. For more information, see [Setting up permissions for configuring metadata tables](#) in the *Amazon S3 User Guide*.

Note

The IAM policy action name is the same for the V1 and V2 API operations.

The following operations are related to GetBucketMetadataConfiguration:

- [CreateBucketMetadataConfiguration](#)
- [DeleteBucketMetadataConfiguration](#)
- [UpdateBucketMetadataInventoryTableConfiguration](#)
- [UpdateBucketMetadataJournalTableConfiguration](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?metadataConfiguration HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

[Bucket](#)

The general purpose bucket that corresponds to the metadata configuration that you want to retrieve.

Required: Yes

[x-amz-expected-bucket-owner](#)

The expected owner of the general purpose bucket that you want to retrieve the metadata table configuration for.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<GetBucketMetadataConfigurationResult>
  <MetadataConfigurationResult>
    <AnnotationTableConfigurationResult>
      <ConfigurationState>string</ConfigurationState>
      <Error>
        <ErrorCode>string</ErrorCode>
        <ErrorMessage>string</ErrorMessage>
      </Error>
      <Role>string</Role>
      <TableArn>string</TableArn>
      <TableName>string</TableName>
```

```

    <TableStatus>string</TableStatus>
  </AnnotationTableConfigurationResult>
  <DestinationResult>
    <TableBucketArn>string</TableBucketArn>
    <TableBucketType>string</TableBucketType>
    <TableNamespace>string</TableNamespace>
  </DestinationResult>
  <InventoryTableConfigurationResult>
    <ConfigurationState>string</ConfigurationState>
    <Error>
      <ErrorCode>string</ErrorCode>
      <ErrorMessage>string</ErrorMessage>
    </Error>
    <TableArn>string</TableArn>
    <TableName>string</TableName>
    <TableStatus>string</TableStatus>
  </InventoryTableConfigurationResult>
  <JournalTableConfigurationResult>
    <Error>
      <ErrorCode>string</ErrorCode>
      <ErrorMessage>string</ErrorMessage>
    </Error>
    <RecordExpiration>
      <Days>integer</Days>
      <Expiration>string</Expiration>
    </RecordExpiration>
    <TableArn>string</TableArn>
    <TableName>string</TableName>
    <TableStatus>string</TableStatus>
  </JournalTableConfigurationResult>
</MetadataConfigurationResult>
</GetBucketMetadataConfigurationResult>

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

GetBucketMetadataConfigurationResult

Root level tag for the GetBucketMetadataConfigurationResult parameters.

Required: Yes

MetadataConfigurationResult

The metadata configuration for a general purpose bucket.

Type: [MetadataConfigurationResult](#) data type

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketMetadataTableConfiguration

Service: Amazon S3

Important

We recommend that you retrieve your S3 Metadata configurations by using the V2 [GetBucketMetadataTableConfiguration](#) API operation. We no longer recommend using the V1 `GetBucketMetadataTableConfiguration` API operation.

If you created your S3 Metadata configuration before July 15, 2025, we recommend that you delete and re-create your configuration by using [CreateBucketMetadataConfiguration](#) so that you can expire journal table records and create a live inventory table.

Retrieves the V1 S3 Metadata configuration for a general purpose bucket. For more information, see [Accelerating data discovery with S3 Metadata](#) in the *Amazon S3 User Guide*.

Note

You can use the V2 `GetBucketMetadataConfiguration` API operation with V1 or V2 metadata table configurations. However, if you try to use the V1 `GetBucketMetadataTableConfiguration` API operation with V2 configurations, you will receive an HTTP 405 Method Not Allowed error.

Make sure that you update your processes to use the new V2 API operations (`CreateBucketMetadataConfiguration`, `GetBucketMetadataConfiguration`, and `DeleteBucketMetadataConfiguration`) instead of the V1 API operations.

Permissions

To use this operation, you must have the `s3:GetBucketMetadataTableConfiguration` permission. For more information, see [Setting up permissions for configuring metadata tables](#) in the *Amazon S3 User Guide*.

The following operations are related to `GetBucketMetadataTableConfiguration`:

- [CreateBucketMetadataTableConfiguration](#)
- [DeleteBucketMetadataTableConfiguration](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?metadataTable HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The general purpose bucket that corresponds to the metadata table configuration that you want to retrieve.

Required: Yes

x-amz-expected-bucket-owner

The expected owner of the general purpose bucket that you want to retrieve the metadata table configuration for.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<GetBucketMetadataTableConfigurationResult>
  <MetadataTableConfigurationResult>
    <S3TablesDestinationResult>
      <TableArn>string</TableArn>
```

```

    <TableBucketArn>string</TableBucketArn>
    <TableName>string</TableName>
    <TableNamespace>string</TableNamespace>
  </S3TablesDestinationResult>
</MetadataTableConfigurationResult>
<Status>string</Status>
<Error>
  <ErrorCode>string</ErrorCode>
  <ErrorMessage>string</ErrorMessage>
</Error>
</GetBucketMetadataTableConfigurationResult>

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

GetBucketMetadataTableConfigurationResult

Root level tag for the GetBucketMetadataTableConfigurationResult parameters.

Required: Yes

Error

If the CreateBucketMetadataTableConfiguration request succeeds, but S3 Metadata was unable to create the table, this structure contains the error code and error message.

Type: [ErrorDetails](#) data type

MetadataTableConfigurationResult

The V1 S3 Metadata configuration for a general purpose bucket.

Type: [MetadataTableConfigurationResult](#) data type

Status

The status of the metadata table. The status values are:

- CREATING - The metadata table is in the process of being created in the specified table bucket.
- ACTIVE - The metadata table has been created successfully, and records are being delivered to the table.

- FAILED - Amazon S3 is unable to create the metadata table, or Amazon S3 is unable to deliver records. See `ErrorDetails` for details.

Type: String

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketMetricsConfiguration

Service: Amazon S3

Gets a metrics configuration (specified by the metrics configuration ID) from the bucket. Note that this doesn't include the daily storage metrics.

Note

Directory buckets - For directory buckets, you must make requests for this API operation to the Regional endpoint. These endpoints support path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name` . Virtual-hosted-style requests aren't supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

Permissions

To use this operation, you must have permissions to perform the `s3:GetMetricsConfiguration` action. The bucket owner has this permission by default. The bucket owner can grant this permission to others. For more information about permissions, see [Permissions Related to Bucket Subresource Operations](#) and [Managing Access Permissions to Your Amazon S3 Resources](#).

- **General purpose bucket permissions** - The `s3:GetMetricsConfiguration` permission is required in a policy. For more information about general purpose buckets permissions, see [Using Bucket Policies and User Policies](#) in the *Amazon S3 User Guide*.
- **Directory bucket permissions** - To grant access to this API operation, you must have the `s3express:GetMetricsConfiguration` permission in an IAM identity-based policy instead of a bucket policy. Cross-account access to this API operation isn't supported. This operation can only be performed by the AWS account that owns the resource. For more information about directory bucket policies and permissions, see [AWS Identity and Access Management \(IAM\) for S3 Express One Zone](#) in the *Amazon S3 User Guide*.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `s3express-control.region-code.amazonaws.com`.

For information about CloudWatch request metrics for Amazon S3, see [Monitoring Metrics with Amazon CloudWatch](#).

The following operations are related to `GetBucketMetricsConfiguration`:

- [PutBucketMetricsConfiguration](#)
- [DeleteBucketMetricsConfiguration](#)
- [ListBucketMetricsConfigurations](#)
- [Monitoring Metrics with Amazon CloudWatch](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?metrics&id=Id HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

[Bucket](#)

The name of the bucket containing the metrics configuration to retrieve.

Directory buckets - When you use this operation with a directory bucket, you must use path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name`. Virtual-hosted-style requests aren't supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must also follow the format `bucket-base-name--zone-id--x-s3` (for example, `DOC-EXAMPLE-BUCKET--usw2-az1--x-s3`). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*.

Required: Yes

id

The ID used to identify the metrics configuration. The ID has a 64 character limit and can only contain letters, numbers, periods, dashes, and underscores.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Note

For directory buckets, this header is not supported in this API operation. If you specify this header, the request fails with the HTTP status code 501 Not Implemented.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<MetricsConfiguration>
  <Id>string</Id>
  <Filter>
    <AccessPointArn>string</AccessPointArn>
    <And>
      <AccessPointArn>string</AccessPointArn>
      <Prefix>string</Prefix>
      <Tag>
        <Key>string</Key>
        <Value>string</Value>
      </Tag>
      ...
    </And>
    <Prefix>string</Prefix>
```

```
<Tag>
  <Key>string</Key>
  <Value>string</Value>
</Tag>
</Filter>
</MetricsConfiguration>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

MetricsConfiguration

Root level tag for the MetricsConfiguration parameters.

Required: Yes

Filter

Specifies a metrics configuration filter. The metrics configuration will only include objects that meet the filter's criteria. A filter must be a prefix, an object tag, an access point ARN, or a conjunction (MetricsAndOperator).

Note

Metrics configurations for directory buckets do not support tag filters.

Type: [MetricsFilter](#) data type

Id

The ID used to identify the metrics configuration. The ID has a 64 character limit and can only contain letters, numbers, periods, dashes, and underscores.

Type: String

Examples

First Sample Request

Retrieve a metrics configuration that filters metrics based on a specified prefix.

```
GET /?metrics&id=Documents HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
x-amz-date: Thu, 15 Nov 2016 00:17:21 GMT
Authorization: signatureValue
```

First Sample Response

This example illustrates one usage of `GetBucketMetricsConfiguration`.

```
HTTP/1.1 200 OK
x-amz-id-2:
ITnGT1y4REXAMPLEPi4hk1TXouTf0hccUjo0iCPEXAMPLEutBj3M7fPGlW02SEWp
x-amz-request-id: 51991EXAMPLE5321
Date: Thu, 15 Nov 2016 00:17:22 GMT
Server: AmazonS3
Content-Length: 180

<?xml version="1.0" encoding="UTF-8"?>
<MetricsConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Id>Documents</Id>
  <Filter>
    <Prefix>documents/</Prefix>
  </Filter>
</MetricsConfiguration>
```

Second Sample Request

Retrieve a metrics configuration that enables metrics for objects that start with a particular prefix and have specific tags applied.

```
GET /?metrics&id=ImportantBlueDocuments HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
x-amz-date: Thu, 15 Nov 2016 00:17:21 GMT
Authorization: signatureValue
```

Second Sample Response

This example illustrates one usage of GetBucketMetricsConfiguration.

```
HTTP/1.1 200 OK
x-amz-id-2:
ITnGT1y4REXAMPLEPi4hk1TXouTf0hccUjo0iCPEXAMPLEutBj3M7fPG1W02SEWp
x-amz-request-id: 51991EXAMPLE5321
Date: Thu, 15 Nov 2016 00:17:22 GMT
Server: AmazonS3
Content-Length: 480

<?xml version="1.0" encoding="UTF-8"?>
<MetricsConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Id>ImportantBlueDocuments</Id>
  <Filter>
    <And>
      <Prefix>documents/</Prefix>
      <Tag>
        <Key>priority</Key>
        <Value>high</Value>
      </Tag>
      <Tag>
        <Key>class</Key>
        <Value>blue</Value>
      </Tag>
    </And>
  </Filter>
</MetricsConfiguration>
```

Third Sample Request

Retrieve a metrics configuration that enables metrics for a specific access point.

```
GET /?metrics&id=ImportantDocumentsAccessPoint HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
x-amz-date: Thu, 26 Aug 2021 00:17:21 GMT
Authorization: signatureValue
```

Third Sample Response

This example illustrates one usage of `GetBucketMetricsConfiguration`.

```
HTTP/1.1 200 OK
x-amz-id-2:
ITnGT1y4REXAMPLEPi4hk1TXouTf0hccUjo0iCPEXAMPLEutBj3M7fPGlW02SEWp
x-amz-request-id: 51991EXAMPLE5321
Date: Thu, 26 Aug 2021 00:17:22 GMT
Server: AmazonS3
Content-Length: 480

<?xml version="1.0" encoding="UTF-8"?>
<MetricsConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Id>ImportantDocumentsAccessPoint</Id>
  <Filter>
    <AccessPointArn>arn:aws:s3:us-west-2:123456789012:accesspoint/test</
AccessPointArn>
  </Filter>
</MetricsConfiguration>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketNotification

Service: Amazon S3

Note

This operation is not supported for directory buckets.

No longer used, see [GetBucketNotificationConfiguration](#).

Request Syntax

```
GET /?notification HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket for which to get the notification configuration.

When you use this API operation with an access point, provide the alias of the access point in place of the bucket name.

When you use this API operation with an Object Lambda access point, provide the alias of the Object Lambda access point in place of the bucket name. If the Object Lambda access point alias in a request is not valid, the error code `InvalidAccessPointAliasError` is returned. For more information about `InvalidAccessPointAliasError`, see [List of Error Codes](#).

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<NotificationConfiguration>
  <TopicConfiguration>
    <Event>string</Event>
    <Event>string</Event>
    ...
    <Id>string</Id>
    <Topic>string</Topic>
  </TopicConfiguration>
  <QueueConfiguration>
    <Event>string</Event>
    <Event>string</Event>
    ...
    <Id>string</Id>
    <Queue>string</Queue>
  </QueueConfiguration>
  <CloudFunctionConfiguration>
    <CloudFunction>string</CloudFunction>
    <Event>string</Event>
    <Event>string</Event>
    ...
    <Id>string</Id>
    <InvocationRole>string</InvocationRole>
  </CloudFunctionConfiguration>
</NotificationConfiguration>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

NotificationConfiguration

Root level tag for the NotificationConfiguration parameters.

Required: Yes

CloudFunctionConfiguration

Container for specifying the AWS Lambda notification configuration.

Type: [CloudFunctionConfiguration](#) data type

QueueConfiguration

This data type is deprecated. This data type specifies the configuration for publishing messages to an Amazon Simple Queue Service (Amazon SQS) queue when Amazon S3 detects specified events.

Type: [QueueConfigurationDeprecated](#) data type

TopicConfiguration

This data type is deprecated. A container for specifying the configuration for publication of messages to an Amazon Simple Notification Service (Amazon SNS) topic when Amazon S3 detects specified events.

Type: [TopicConfigurationDeprecated](#) data type

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketNotificationConfiguration

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Returns the notification configuration of a bucket.

If notifications are not enabled on the bucket, the action returns an empty `NotificationConfiguration` element.

By default, you must be the bucket owner to read the notification configuration of a bucket. However, the bucket owner can use a bucket policy to grant permission to other users to read this configuration with the `s3:GetBucketNotification` permission.

When you use this API operation with an access point, provide the alias of the access point in place of the bucket name.

When you use this API operation with an Object Lambda access point, provide the alias of the Object Lambda access point in place of the bucket name. If the Object Lambda access point alias in a request is not valid, the error code `InvalidAccessPointAliasError` is returned. For more information about `InvalidAccessPointAliasError`, see [List of Error Codes](#).

For more information about setting and reading the notification configuration on a bucket, see [Setting Up Notification of Bucket Events](#). For more information about bucket policies, see [Using Bucket Policies](#).

The following action is related to `GetBucketNotification`:

- [PutBucketNotification](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?notification HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket for which to get the notification configuration.

When you use this API operation with an access point, provide the alias of the access point in place of the bucket name.

When you use this API operation with an Object Lambda access point, provide the alias of the Object Lambda access point in place of the bucket name. If the Object Lambda access point alias in a request is not valid, the error code `InvalidAccessPointAliasError` is returned. For more information about `InvalidAccessPointAliasError`, see [List of Error Codes](#).

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<NotificationConfiguration>
  <TopicConfiguration>
    <Event>string</Event>
    ...
```

```

<Filter>
  <S3Key>
    <FilterRule>
      <Name>string</Name>
      <Value>string</Value>
    </FilterRule>
    ...
  </S3Key>
</Filter>
<Id>string</Id>
<Topic>string</Topic>
</TopicConfiguration>
...
<QueueConfiguration>
  <Event>string</Event>
  ...
  <Filter>
    <S3Key>
      <FilterRule>
        <Name>string</Name>
        <Value>string</Value>
      </FilterRule>
      ...
    </S3Key>
  </Filter>
  <Id>string</Id>
  <Queue>string</Queue>
</QueueConfiguration>
...
<CloudFunctionConfiguration>
  <Event>string</Event>
  ...
  <Filter>
    <S3Key>
      <FilterRule>
        <Name>string</Name>
        <Value>string</Value>
      </FilterRule>
      ...
    </S3Key>
  </Filter>
  <Id>string</Id>
  <CloudFunction>string</CloudFunction>
</CloudFunctionConfiguration>

```



```
...  
<EventBridgeConfiguration>  
</EventBridgeConfiguration>  
</NotificationConfiguration>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[NotificationConfiguration](#)

Root level tag for the NotificationConfiguration parameters.

Required: Yes

[CloudFunctionConfiguration](#)

Describes the AWS Lambda functions to invoke and the events for which to invoke them.

Type: Array of [LambdaFunctionConfiguration](#) data types

[EventBridgeConfiguration](#)

Enables delivery of events to Amazon EventBridge.

Type: [EventBridgeConfiguration](#) data type

[QueueConfiguration](#)

The Amazon Simple Queue Service queues to publish messages to and the events for which to publish messages.

Type: Array of [QueueConfiguration](#) data types

[TopicConfiguration](#)

The topic to which notifications are sent and the events for which notifications are generated.

Type: Array of [TopicConfiguration](#) data types

Examples

Sample Request

This request returns the notification configuration on the bucket `amzn-s3-demo-bucket.s3.<Region>.amazonaws.com`.

```
GET ?notification HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Wed, 15 Oct 2014 16:59:03 GMT
Authorization: authorization string
```

Sample Response

This response returns that the notification configuration for the specified bucket.

```
HTTP/1.1 200 OK
x-amz-id-2: YgIPIfBiKa2bj0KMgUAdQkf3ShJT00pXUueF6QKo
x-amz-request-id: 236A8905248E5A02
Date: Wed, 15 Oct 2014 16:59:04 GMT
Server: AmazonS3
<?xml version="1.0" encoding="UTF-8"?>

<NotificationConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <TopicConfiguration>
    <Id>YjVkJM2Y0YmUtNGI3NC00ZjQyLWEwNGItNDIyYWUxY2I0N2M4</Id>
    <Topic>arn:aws:sns:us-east-1:account-id:s3notificationtopic2</Topic>
    <Event>s3:ReducedRedundancyLostObject</Event>
    <Event>s3:ObjectCreated:*</Event>
  </TopicConfiguration>
</NotificationConfiguration>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)

- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketOwnershipControls

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Retrieves OwnershipControls for an Amazon S3 bucket. To use this operation, you must have the `s3:GetBucketOwnershipControls` permission. For more information about Amazon S3 permissions, see [Specifying permissions in a policy](#).

Note

A bucket doesn't have OwnershipControls settings in the following cases:

- The bucket was created before the BucketOwnerEnforced ownership setting was introduced and you've never explicitly applied this value
- You've manually deleted the bucket ownership control value using the DeleteBucketOwnershipControls API operation.

By default, Amazon S3 sets OwnershipControls for all newly created buckets.

For information about Amazon S3 Object Ownership, see [Using Object Ownership](#).

The following operations are related to GetBucketOwnershipControls:

- [PutBucketOwnershipControls](#)
- [DeleteBucketOwnershipControls](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?ownershipControls HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the Amazon S3 bucket whose OwnershipControls you want to retrieve.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<OwnershipControls>
  <Rule>
    <ObjectOwnership>string</ObjectOwnership>
  </Rule>
  ...
</OwnershipControls>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

OwnershipControls

Root level tag for the OwnershipControls parameters.

Required: Yes

Rule

The container element for an ownership control rule.

Type: Array of [OwnershipControlsRule](#) data types

Examples

Sample GetBucketOwnershipControls Request for BucketOwnerEnforced

This example illustrates one usage of GetBucketOwnershipControls.

```
GET /amzn-s3-demo-bucket?/ownershipControls HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Mon, 29 Nov 2021 00:17:22 GMT
Authorization: signatureValue;
```

Sample GetBucketOwnershipControls Response

This example illustrates one usage of GetBucketOwnershipControls.

```
HTTP/1.1 200 OK
x-amz-id-2: Adphn7MaAHDEg9mh5JmcTN8mzyVX0JhIztSiQNaqTxnXXcYi4uiZbYdwWC3JXmh/
XXVUUQw04Vs=
x-amz-request-id: 252631E05F84A415
Date: Mon, 29 Nov 2021 00:17:22 GMT
Server: AmazonS3
Content-Length: 194

<OwnershipControls xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Rule>
```

```
<ObjectOwnership>BucketOwnerEnforced</ObjectOwnership>
</Rule>
</OwnershipControls>
```

Sample GetBucketOwnershipControls Request for BucketOwnerPreferred

This example illustrates one usage of GetBucketOwnershipControls.

```
GET /amzn-s3-demo-bucket?/ownershipControls HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Thu, 18 Jun 2017 00:17:22 GMT
Authorization: signatureValue;
```

Sample GetBucketOwnershipControls Response

This example illustrates one usage of GetBucketOwnershipControls.

```
HTTP/1.1 200 OK
x-amz-id-2: Adphn7MaAHDEg9mh5JmcTN8mzyVX0JhIztSiQNaqTxnXXcYi4uiZbYdwWC3JXmh/
XXVUUQw04Vs=
x-amz-request-id: 252631E05F84A415
Date: Thu, 18 Jun 2020 00:17:22 GMT
Server: AmazonS3
Content-Length: 194

<OwnershipControls xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Rule>
    <ObjectOwnership>BucketOwnerPreferred</ObjectOwnership>
  </Rule>
</OwnershipControls>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)

- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketPolicy

Service: Amazon S3

Returns the policy of a specified bucket.

Note

Directory buckets - For directory buckets, you must make requests for this API operation to the Regional endpoint. These endpoints support path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name` . Virtual-hosted-style requests aren't supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

Permissions

If you are using an identity other than the root user of the AWS account that owns the bucket, the calling identity must both have the `GetBucketPolicy` permissions on the specified bucket and belong to the bucket owner's account in order to use this operation.

If you don't have `GetBucketPolicy` permissions, Amazon S3 returns a `403 Access Denied` error. If you have the correct permissions, but you're not using an identity that belongs to the bucket owner's account, Amazon S3 returns a `405 Method Not Allowed` error.

Important

To ensure that bucket owners don't inadvertently lock themselves out of their own buckets, the root principal in a bucket owner's AWS account can perform the `GetBucketPolicy`, `PutBucketPolicy`, and `DeleteBucketPolicy` API actions, even if their bucket policy explicitly denies the root principal's access. Bucket owner root principals can only be blocked from performing these API actions by VPC endpoint policies and AWS Organizations policies.

- **General purpose bucket permissions** - The `s3:GetBucketPolicy` permission is required in a policy. For more information about general purpose buckets bucket policies, see [Using Bucket Policies and User Policies](#) in the *Amazon S3 User Guide*.

- **Directory bucket permissions** - To grant access to this API operation, you must have the `s3express:GetBucketPolicy` permission in an IAM identity-based policy instead of a bucket policy. Cross-account access to this API operation isn't supported. This operation can only be performed by the AWS account that owns the resource. For more information about directory bucket policies and permissions, see [AWS Identity and Access Management \(IAM\) for S3 Express One Zone](#) in the *Amazon S3 User Guide*.

Example bucket policies

General purpose buckets example bucket policies - See [Bucket policy examples](#) in the *Amazon S3 User Guide*.

Directory bucket example bucket policies - See [Example bucket policies for S3 Express One Zone](#) in the *Amazon S3 User Guide*.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `s3express-control.region-code.amazonaws.com`.

The following action is related to `GetBucketPolicy`:

- [GetObject](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?policy HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name to get the bucket policy for.

Directory buckets - When you use this operation with a directory bucket, you must use path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name`. Virtual-hosted-style requests aren't supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must also follow the format `bucket-base-name--zone-id--x-s3` (for example, `DOC-EXAMPLE-BUCKET--usw2-az1--x-s3`). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*.

Access points - When you use this API operation with an access point, provide the alias of the access point in place of the bucket name.

Object Lambda access points - When you use this API operation with an Object Lambda access point, provide the alias of the Object Lambda access point in place of the bucket name. If the Object Lambda access point alias in a request is not valid, the error code `InvalidAccessPointAliasError` is returned. For more information about `InvalidAccessPointAliasError`, see [List of Error Codes](#).

Note

Object Lambda access points are not supported by directory buckets.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code `403 Forbidden` (access denied).

Note

For directory buckets, this header is not supported in this API operation. If you specify this header, the request fails with the HTTP status code `501 Not Implemented`.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200  
  
{ Policy in JSON format }
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[Policy](#)

Examples

Sample Request for general purpose buckets

The following request returns the policy of the specified bucket.

```
GET ?policy HTTP/1.1  
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com  
Date: Wed, 28 Oct 2009 22:32:00 GMT  
Authorization: authorization string
```

Sample Response for general purpose buckets

This example illustrates one usage of GetBucketPolicy.

```
HTTP/1.1 200 OK  
x-amz-id-2: Uuag1LuByru9p04SAMPLEAtRPfTa0Fg==  
x-amz-request-id: 656c76696e67SAMPLE57374  
Date: Tue, 04 Apr 2010 20:34:56 GMT  
Connection: keep-alive
```

Server: AmazonS3

```
{
  "Version": "2008-10-17",
  "Id": "aaaa-bbbb-cccc-dddd",
  "Statement" : [
    {
      "Effect": "Deny",
      "Sid": "1",
      "Principal" : {
        "AWS": ["111122223333", "444455556666"]
      },
      "Action": ["s3:*"],
      "Resource": "arn:aws:s3:::bucket/*"
    }
  ]
}
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketPolicyStatus

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Retrieves the policy status for an Amazon S3 bucket, indicating whether the bucket is public. In order to use this operation, you must have the `s3:GetBucketPolicyStatus` permission. For more information about Amazon S3 permissions, see [Specifying Permissions in a Policy](#).

For more information about when Amazon S3 considers a bucket public, see [The Meaning of "Public"](#).

The following operations are related to `GetBucketPolicyStatus`:

- [Using Amazon S3 Block Public Access](#)
- [GetPublicAccessBlock](#)
- [PutPublicAccessBlock](#)
- [DeletePublicAccessBlock](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?policyStatus HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the Amazon S3 bucket whose policy status you want to retrieve.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<PolicyStatus>
  <IsPublic>boolean</IsPublic>
</PolicyStatus>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

PolicyStatus

Root level tag for the PolicyStatus parameters.

Required: Yes

IsPublic

The policy status for this bucket. TRUE indicates that this bucket is public. FALSE indicates that the bucket is not public.

Type: Boolean

Examples

Sample Request

The following request gets a bucket policy status.

```
GET /<bucket-name>?policyStatus HTTP/1.1
Host: <bucket-name>.s3.<Region>.amazonaws.com
x-amz-date: <Thu, 15 Nov 2016 00:17:21 GMT>
Authorization: <signatureValue>
```

Sample Response

This example illustrates one usage of `GetBucketPolicyStatus`.

```
HTTP/1.1 200 OK
x-amz-id-2:
ITnGT1y4REXAMPLEPi4hk1TXouTf0hccUjo0iCPEXAMPLEutBj3M7fPG1W02SEWp
x-amz-request-id: 51991EXAMPLE5321
Date: Thu, 15 Nov 2016 00:17:22 GMT
Server: AmazonS3
Content-Length: 0

<PolicyStatus>
  <IsPublic>TRUE</IsPublic>
</PolicyStatus>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketReplication

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Returns the replication configuration of a bucket.

Note

It can take a while to propagate the put or delete a replication configuration to all Amazon S3 systems. Therefore, a get request soon after put or delete can return a wrong result.

For information about replication configuration, see [Replication](#) in the *Amazon S3 User Guide*.

This action requires permissions for the `s3:GetReplicationConfiguration` action. For more information about permissions, see [Using Bucket Policies and User Policies](#).

If you include the `Filter` element in a replication configuration, you must also include the `DeleteMarkerReplication` and `Priority` elements. The response also returns those elements.

For information about `GetBucketReplication` errors, see [List of replication-related error codes](#)

The following operations are related to `GetBucketReplication`:

- [PutBucketReplication](#)
- [DeleteBucketReplication](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?replication HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name for which to get the replication information.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<ReplicationConfiguration>
  <Role>string</Role>
  <Rule>
    <DeleteMarkerReplication>
      <Status>string</Status>
    </DeleteMarkerReplication>
    <Destination>
      <AccessControlTranslation>
        <Owner>string</Owner>
      </AccessControlTranslation>
      <Account>string</Account>
      <Bucket>string</Bucket>
      <EncryptionConfiguration>
```

```

    <ReplicaKmsKeyID>string</ReplicaKmsKeyID>
  </EncryptionConfiguration>
  <Metrics>
    <EventThreshold>
      <Minutes>integer</Minutes>
    </EventThreshold>
    <Status>string</Status>
  </Metrics>
  <ReplicationTime>
    <Status>string</Status>
    <Time>
      <Minutes>integer</Minutes>
    </Time>
  </ReplicationTime>
  <StorageClass>string</StorageClass>
</Destination>
<ExistingObjectReplication>
  <Status>string</Status>
</ExistingObjectReplication>
<Filter>
  <And>
    <Prefix>string</Prefix>
    <Tag>
      <Key>string</Key>
      <Value>string</Value>
    </Tag>
    ...
  </And>
  <Prefix>string</Prefix>
  <Tag>
    <Key>string</Key>
    <Value>string</Value>
  </Tag>
</Filter>
<ID>string</ID>
<Prefix>string</Prefix>
<Priority>integer</Priority>
<SourceSelectionCriteria>
  <ReplicaModifications>
    <Status>string</Status>
  </ReplicaModifications>
  <SseKmsEncryptedObjects>
    <Status>string</Status>
  </SseKmsEncryptedObjects>

```

```
</SourceSelectionCriteria>
<Status>string</Status>
</Rule>
...
</ReplicationConfiguration>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

ReplicationConfiguration

Root level tag for the ReplicationConfiguration parameters.

Required: Yes

Role

The Amazon Resource Name (ARN) of the AWS Identity and Access Management (IAM) role that Amazon S3 assumes when replicating objects. For more information, see [How to Set Up Replication](#) in the *Amazon S3 User Guide*.

Type: String

Rule

A container for one or more replication rules. A replication configuration must have at least one rule and can contain a maximum of 1,000 rules.

Type: Array of [ReplicationRule](#) data types

Examples

Sample Request: Retrieve replication configuration information

The following GET request retrieves information about the replication configuration set for the amzn-s3-demo-bucket bucket:

```
GET /?replication HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
```

Date: Tue, 10 Feb 2015 00:17:21 GMT
 Authorization: authorization string

Sample Response

The following response shows that replication is enabled on the bucket. The empty prefix indicates that Amazon S3 will replicate all objects that are created in the `amzn-s3-demo-bucket` bucket. The `Destination` element identifies the target bucket where Amazon S3 creates the object replicas, and the storage class (`STANDARD_IA`) that Amazon S3 uses when creating replicas.

Amazon S3 assumes the specified IAM role to replicate objects on behalf of the bucket owner, which is the AWS account that created the bucket.

```

HTTP/1.1 200 OK
x-amz-id-2:
ITnGT1y4RyTmXa3rPi4hklTXouTf0hccUjo0iCPjz6FnfIutBj3M7fPGlW02SEWp
x-amz-request-id: 51991C342example
Date: Tue, 10 Feb 2015 00:17:23 GMT
Server: AmazonS3
Content-Length: contentlength

<?xml version="1.0" encoding="UTF-8"?>
<ReplicationConfiguration>
  <Role>arn:aws:iam::35667example:role/CrossRegionReplicationRoleForS3</
Role>
  <Rule>
    <ID>rule1</ID>
    <Status>Enabled</Status>
    <Priority>1</Priority>
    <DeleteMarkerReplication>
      <Status>Disabled</Status>
    </DeleteMarkerReplication>
    <Filter>
      <And>
        <Prefix>TaxDocs</Prefix>
        <Tag>
          <Key>key1</Key>
          <Value>value1</Value>
        </Tag>
        <Tag>
          <Key>key1</Key>

```

```
        <Value>value1</Value>
      </Tag>
    </And>
  </Filter>
<Destination>
  <Bucket>arn:aws:s3:::exampletargetbucket</Bucket>
</Destination>
</Rule>
</ReplicationConfiguration>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketRequestPayment

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Returns the request payment configuration of a bucket. To use this version of the operation, you must be the bucket owner. For more information, see [Requester Pays Buckets](#).

The following operations are related to GetBucketRequestPayment:

- [ListObjects](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?requestPayment HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket for which to get the payment request configuration

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<RequestPaymentConfiguration>
  <Payer>string</Payer>
</RequestPaymentConfiguration>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

RequestPaymentConfiguration

Root level tag for the RequestPaymentConfiguration parameters.

Required: Yes

Payer

Specifies who pays for the download and request fees.

Type: String

Valid Values: Requester | BucketOwner

Examples

Sample Request

The following request returns the payer for the bucket, amzn-s3-demo-bucket.

```
GET ?requestPayment HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Wed, 01 Mar 2009 12:00:00 GMT
Authorization: authorization string
```

Sample Response

This response shows that the bucket is a Requester Pays bucket, meaning the person requesting a download from this bucket pays the transfer fees.

```
HTTP/1.1 200 OK
x-amz-id-2:
YgIPIfBiKa2bj0KMg95r/0zo3emzU4dzsD4rcKCHQUAdQkf3ShJT00pXUueF6QKo
x-amz-request-id: 236A8905248E5A01
Date: Wed, 01 Mar 2009 12:00:00 GMT
Content-Type: [type]
Content-Length: 0
Connection: close
Server: AmazonS3

<?xml version="1.0" encoding="UTF-8"?>
<RequestPaymentConfiguration xmlns="http://s3.amazonaws.com/
doc/2006-03-01/">
  <Payer>Requester</Payer>
</RequestPaymentConfiguration>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)

- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketTagging

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Returns the tag set associated with the general purpose bucket.

To use this operation, you must have permission to perform the `s3:GetBucketTagging` action. By default, the bucket owner has this permission and can grant this permission to others.

GetBucketTagging has the following special error:

- Error code: `NoSuchTagSet`
 - Description: There is no tag set associated with the bucket.

The following operations are related to GetBucketTagging:

- [PutBucketTagging](#)
- [DeleteBucketTagging](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?tagging HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

[Bucket](#)

The name of the bucket for which to get the tagging information.

Required: Yes

[x-amz-expected-bucket-owner](#)

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<Tagging>
  <TagSet>
    <Tag>
      <Key>string</Key>
      <Value>string</Value>
    </Tag>
  </TagSet>
</Tagging>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[Tagging](#)

Root level tag for the Tagging parameters.

Required: Yes

[TagSet](#)

Contains the tag set.

Type: Array of [Tag](#) data types

Examples

Sample Request

The following request returns the tag set of the specified bucket.

```
GET ?tagging HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Wed, 28 Oct 2009 22:32:00 GMT
Authorization: authorization string
```

Sample Response

Delete the metric configuration with a specified ID, which disables the CloudWatch metrics with the ExampleMetrics value for the FilterId dimension.

```
HTTP/1.1 200 OK
Date: Wed, 25 Nov 2009 12:00:00 GMT
Connection: close
Server: AmazonS3

<Tagging>
  <TagSet>
    <Tag>
      <Key>Project</Key>
      <Value>Project One</Value>
    </Tag>
    <Tag>
      <Key>User</Key>
      <Value>jsmith</Value>
    </Tag>
```

```
</TagSet>  
</Tagging>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketVersioning

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Returns the versioning state of a bucket.

To retrieve the versioning state of a bucket, you must be the bucket owner.

This implementation also returns the MFA Delete status of the versioning state. If the MFA Delete status is enabled, the bucket owner must use an authentication device to change the versioning state of the bucket.

The following operations are related to GetBucketVersioning:

- [GetObject](#)
- [PutObject](#)
- [DeleteObject](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?versioning HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket for which to get the versioning information.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<VersioningConfiguration>
  <Status>string</Status>
  <MfaDelete>string</MfaDelete>
</VersioningConfiguration>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

VersioningConfiguration

Root level tag for the VersioningConfiguration parameters.

Required: Yes

MFADelete

Specifies whether MFA delete is enabled in the bucket versioning configuration. This element is only returned if the bucket has been configured with MFA delete. If the bucket has never been so configured, this element is not returned.

Type: String

Valid Values: Enabled | Disabled

Status

The versioning state of the bucket.

Type: String

Valid Values: Enabled | Suspended

Examples

Example

This example returns the versioning state of amzn-s3-demo-bucket.

```
GET /?versioning HTTP/1.1
Host:amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Wed, 12 Oct 2009 17:50:00 GMT
Authorization: authorization string
Content-Type: text/plain
```

Example

There are three versioning states:

If you enabled versioning on a bucket, the response is:

```
<VersioningConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Status>Enabled</Status>
</VersioningConfiguration>
```

Example

If you suspended versioning on a bucket, the response is:

```
<VersioningConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Status>Suspended</Status>
</VersioningConfiguration>
```

Example

If you never enabled (or suspended) versioning on a bucket, the response is:

```
<VersioningConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketWebsite

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Returns the website configuration for a bucket. To host website on Amazon S3, you can configure a bucket as website by adding a website configuration. For more information about hosting websites, see [Hosting Websites on Amazon S3](#).

This GET action requires the `S3:GetBucketWebsite` permission. By default, only the bucket owner can read the bucket website configuration. However, bucket owners can allow other users to read the website configuration by writing a bucket policy granting them the `S3:GetBucketWebsite` permission.

The following operations are related to `GetBucketWebsite`:

- [DeleteBucketWebsite](#)
- [PutBucketWebsite](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?website HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name for which to get the website configuration.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<WebsiteConfiguration>
  <RedirectAllRequestsTo>
    <HostName>string</HostName>
    <Protocol>string</Protocol>
  </RedirectAllRequestsTo>
  <IndexDocument>
    <Suffix>string</Suffix>
  </IndexDocument>
  <ErrorDocument>
    <Key>string</Key>
  </ErrorDocument>
  <RoutingRules>
    <RoutingRule>
      <Condition>
        <HttpErrorCodeReturnedEquals>string</HttpErrorCodeReturnedEquals>
        <KeyPrefixEquals>string</KeyPrefixEquals>
      </Condition>
      <Redirect>
        <HostName>string</HostName>
        <HttpRedirectCode>string</HttpRedirectCode>
        <Protocol>string</Protocol>
        <ReplaceKeyPrefixWith>string</ReplaceKeyPrefixWith>
        <ReplaceKeyWith>string</ReplaceKeyWith>
      </Redirect>
    </RoutingRule>
  </RoutingRules>
</WebsiteConfiguration>
```

```
</Redirect>
</RoutingRule>
</RoutingRules>
</WebsiteConfiguration>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

WebsiteConfiguration

Root level tag for the WebsiteConfiguration parameters.

Required: Yes

ErrorDocument

The object key name of the website error document to use for 4XX class errors.

Type: [ErrorDocument](#) data type

IndexDocument

The name of the index document for the website (for example `index.html`).

Type: [IndexDocument](#) data type

RedirectAllRequestsTo

Specifies the redirect behavior of all requests to a website endpoint of an Amazon S3 bucket.

Type: [RedirectAllRequestsTo](#) data type

RoutingRules

Rules that define when a redirect is applied and the redirect behavior.

Type: Array of [RoutingRule](#) data types

Examples

Sample Request

This request retrieves website configuration on the specified bucket.

```
GET ?website HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Thu, 27 Jan 2011 00:49:20 GMT
Authorization: AWS AKIAIOSFODNN7EXAMPLE:n0Nhek72Ufg/u7Sm5C1dqRLs8XX=
```

Sample Response

This example illustrates one usage of GetBucketWebsite.

```
HTTP/1.1 200 OK
x-amz-id-2: YgIPIfBiKa2bj0KMgUAdQkf3ShJT00pXUueF6QKo
x-amz-request-id: 3848CD259D811111
Date: Thu, 27 Jan 2011 00:49:26 GMT
Content-Length: 240
Content-Type: application/xml
Transfer-Encoding: chunked
Server: AmazonS3

<?xml version="1.0" encoding="UTF-8"?>
<WebsiteConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <IndexDocument>
    <Suffix>index.html</Suffix>
  </IndexDocument>
  <ErrorDocument>
    <Key>404.html</Key>
  </ErrorDocument>
</WebsiteConfiguration>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)

- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetObject

Service: Amazon S3

Retrieves an object from Amazon S3.

In the `GetObject` request, specify the full key name for the object.

General purpose buckets - Both the virtual-hosted-style requests and the path-style requests are supported. For a virtual hosted-style request example, if you have the object `photos/2006/February/sample.jpg`, specify the object key name as `/photos/2006/February/sample.jpg`. For a path-style request example, if you have the object `photos/2006/February/sample.jpg` in the bucket named `examplebucket`, specify the object key name as `/examplebucket/photos/2006/February/sample.jpg`. For more information about request types, see [HTTP Host Header Bucket Specification](#) in the *Amazon S3 User Guide*.

Directory buckets - Only virtual-hosted-style requests are supported. For a virtual hosted-style request example, if you have the object `photos/2006/February/sample.jpg` in the bucket named `amzn-s3-demo-bucket--usw2-az1--x-s3`, specify the object key name as `/photos/2006/February/sample.jpg`. Also, when you make requests to this API operation, your requests are sent to the Zonal endpoint. These endpoints support virtual-hosted-style requests in the format `https://bucket-name.s3express-zone-id.region-code.amazonaws.com/key-name`. Path-style requests are not supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

Permissions

- **General purpose bucket permissions** - You must have the required permissions in a policy. To use `GetObject`, you must have the `READ` access to the object (or version). If you grant `READ` access to the anonymous user, the `GetObject` operation returns the object without using an authorization header. For more information, see [Specifying permissions in a policy](#) in the *Amazon S3 User Guide*.

If you include a `versionId` in your request header, you must have the `s3:GetObjectVersion` permission to access a specific version of an object. The `s3:GetObject` permission is not required in this scenario.

If you request the current version of an object without a specific `versionId` in the request header, only the `s3:GetObject` permission is required. The `s3:GetObjectVersion` permission is not required in this scenario.

If the object that you request doesn't exist, the error that Amazon S3 returns depends on whether you also have the `s3:ListBucket` permission.

- If you have the `s3:ListBucket` permission on the bucket, Amazon S3 returns an HTTP status code `404 Not Found` error.
- If you don't have the `s3:ListBucket` permission, Amazon S3 returns an HTTP status code `403 Access Denied` error.
- **Directory bucket permissions** - To grant access to this API operation on a directory bucket, we recommend that you use the [CreateSession](#) API operation for session-based authorization. Specifically, you grant the `s3express:CreateSession` permission to the directory bucket in a bucket policy or an IAM identity-based policy. Then, you make the `CreateSession` API call on the bucket to obtain a session token. With the session token in your request header, you can make API requests to this operation. After the session token expires, you make another `CreateSession` API call to generate a new session token for use. AWS CLI or SDKs create session and refresh the session token automatically to avoid service interruptions when a session expires. For more information about authorization, see [CreateSession](#).

If the object is encrypted using SSE-KMS, you must also have the `kms:GenerateDataKey` and `kms:Decrypt` permissions in IAM identity-based policies and AWS KMS key policies for the AWS KMS key.

Storage classes

If the object you are retrieving is stored in the S3 Glacier Flexible Retrieval storage class, the S3 Glacier Deep Archive storage class, the S3 Intelligent-Tiering Archive Access tier, or the S3 Intelligent-Tiering Deep Archive Access tier, before you can retrieve the object you must first restore a copy using [RestoreObject](#). Otherwise, this operation returns an `InvalidObjectState` error. For information about restoring archived objects, see [Restoring Archived Objects](#) in the *Amazon S3 User Guide*.

Directory buckets - Directory buckets only support `EXPRESS_ONEZONE` (the S3 Express One Zone storage class) in Availability Zones and `ONEZONE_IA` (the S3 One Zone-Infrequent

Access storage class) in Dedicated Local Zones. Unsupported storage class values won't write a destination object and will respond with the HTTP status code 400 Bad Request.

Encryption

Encryption request headers, like `x-amz-server-side-encryption`, should not be sent for the `GetObject` requests, if your object uses server-side encryption with Amazon S3 managed encryption keys (SSE-S3), server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS), or dual-layer server-side encryption with AWS KMS keys (DSSE-KMS). If you include the header in your `GetObject` requests for the object that uses these types of keys, you'll get an HTTP 400 Bad Request error.

Directory buckets - For directory buckets, there are only two supported options for server-side encryption: SSE-S3 and SSE-KMS. SSE-C isn't supported. For more information, see [Protecting data with server-side encryption](#) in the *Amazon S3 User Guide*.

Overriding response header values through the request

There are times when you want to override certain response header values of a `GetObject` response. For example, you might override the `Content-Disposition` response header value through your `GetObject` request.

You can override values for a set of response headers. These modified response header values are included only in a successful response, that is, when the HTTP status code 200 OK is returned. The headers you can override using the following query parameters in the request are a subset of the headers that Amazon S3 accepts when you create an object.

The response headers that you can override for the `GetObject` response are `Cache-Control`, `Content-Disposition`, `Content-Encoding`, `Content-Language`, `Content-Type`, and `Expires`.

To override values for a set of response headers in the `GetObject` response, you can use the following query parameters in the request.

- `response-cache-control`
- `response-content-disposition`
- `response-content-encoding`
- `response-content-language`
- `response-content-type`
- `response-expires`

Note

When you use these parameters, you must sign the request by using either an Authorization header or a presigned URL. These parameters cannot be used with an unsigned (anonymous) request.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `Bucket-name.s3express-zone-id.region-code.amazonaws.com`.

The following operations are related to `GetObject`:

- [ListBuckets](#)
- [GetObjectAcl](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /Key+?partNumber=PartNumber&response-cache-control=ResponseCacheControl&response-content-disposition=ResponseContentDisposition&response-content-encoding=ResponseContentEncoding&response-content-language=ResponseContentLanguage&response-content-type=ResponseContentType&response-expires=ResponseExpires&versionId=VersionId HTTP/1.1
```

Host: `Bucket.s3.amazonaws.com`

If-Match: `IfMatch`

If-Modified-Since: `IfModifiedSince`

If-None-Match: `IfNoneMatch`

If-Unmodified-Since: `IfUnmodifiedSince`

Range: `Range`

x-amz-server-side-encryption-customer-algorithm: `SSECustomerAlgorithm`

x-amz-server-side-encryption-customer-key: `SSECustomerKey`

```
x-amz-server-side-encryption-customer-key-MD5: SSECustomerKeyMD5  
x-amz-request-payer: RequestPayer  
x-amz-expected-bucket-owner: ExpectedBucketOwner  
x-amz-checksum-mode: ChecksumMode
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name containing the object.

Directory buckets - When you use this operation with a directory bucket, you must use virtual-hosted-style requests in the format *Bucket-name.s3express-zone-id.region-code.amazonaws.com*. Path-style requests are not supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must follow the format *bucket-base-name--zone-id--x-s3* (for example, *amzn-s3-demo-bucket--usw2-az1--x-s3*). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form *AccessPointName-AccountId.s3-accesspoint.Region.amazonaws.com*. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

Object Lambda access points - When you use this action with an Object Lambda access point, you must direct requests to the Object Lambda access point hostname. The Object Lambda access point hostname takes the form *AccessPointName-AccountId.s3-object-lambda.Region.amazonaws.com*.

Note

Object Lambda access points are not supported by directory buckets.

S3 on Outposts - When you use this action with S3 on Outposts, you must direct requests to the S3 on Outposts hostname. The S3 on Outposts hostname takes the form *AccessPointName-AccountId.outpostID.s3-outposts.Region.amazonaws.com*. When you use this action with S3 on Outposts, the destination bucket must be the Outposts access point ARN or the access point alias. For more information about S3 on Outposts, see [What is S3 on Outposts?](#) in the *Amazon S3 User Guide*.

Required: Yes

If-Match

Return the object only if its entity tag (ETag) is the same as the one specified in this header; otherwise, return a 412 `Precondition Failed` error.

If both of the `If-Match` and `If-Unmodified-Since` headers are present in the request as follows: `If-Match` condition evaluates to `true`, and; `If-Unmodified-Since` condition evaluates to `false`; then, S3 returns 200 `OK` and the data requested.

For more information about conditional requests, see [RFC 7232](#).

If-Modified-Since

Return the object only if it has been modified since the specified time; otherwise, return a 304 `Not Modified` error.

If both of the `If-None-Match` and `If-Modified-Since` headers are present in the request as follows: `If-None-Match` condition evaluates to `false`, and; `If-Modified-Since` condition evaluates to `true`; then, S3 returns 304 `Not Modified` status code.

For more information about conditional requests, see [RFC 7232](#).

If-None-Match

Return the object only if its entity tag (ETag) is different from the one specified in this header; otherwise, return a 304 `Not Modified` error.

If both of the `If-None-Match` and `If-Modified-Since` headers are present in the request as follows: `If-None-Match` condition evaluates to `false`, and; `If-Modified-Since` condition evaluates to `true`; then, S3 returns 304 `Not Modified` HTTP status code.

For more information about conditional requests, see [RFC 7232](#).

If-Unmodified-Since

Return the object only if it has not been modified since the specified time; otherwise, return a 412 Precondition Failed error.

If both of the If-Match and If-Unmodified-Since headers are present in the request as follows: If-Match condition evaluates to true, and; If-Unmodified-Since condition evaluates to false; then, S3 returns 200 OK and the data requested.

For more information about conditional requests, see [RFC 7232](#).

Key

Key of the object to get.

Length Constraints: Minimum length of 1.

Required: Yes

partNumber

Part number of the object being read. This is a positive integer between 1 and 10,000.

Effectively performs a 'ranged' GET request for the part specified. Useful for downloading just a part of an object.

Range

Downloads the specified byte range of an object. For more information about the HTTP Range header, see <https://www.rfc-editor.org/rfc/rfc9110.html#name-range>.

Note

Amazon S3 doesn't support retrieving multiple ranges of data per GET request.

response-cache-control

Sets the Cache-Control header of the response.

response-content-disposition

Sets the Content-Disposition header of the response.

response-content-encoding

Sets the Content-Encoding header of the response.

response-content-language

Sets the Content-Language header of the response.

response-content-type

Sets the Content-Type header of the response.

response-expires

Sets the Expires header of the response.

versionId

Version ID used to reference a specific version of the object.

By default, the `GetObject` operation returns the current version of an object. To return a different version, use the `versionId` subresource.

Note

- If you include a `versionId` in your request header, you must have the `s3:GetObjectVersion` permission to access a specific version of an object. The `s3:GetObject` permission is not required in this scenario.
- If you request the current version of an object without a specific `versionId` in the request header, only the `s3:GetObject` permission is required. The `s3:GetObjectVersion` permission is not required in this scenario.
- **Directory buckets** - S3 Versioning isn't enabled and supported for directory buckets. For this API operation, only the `null` value of the version ID is supported by directory buckets. You can only specify `null` to the `versionId` query parameter in the request.

For more information about versioning, see [PutBucketVersioning](#).

x-amz-checksum-mode

To retrieve the checksum, this mode must be enabled.

Valid Values: ENABLED

[x-amz-expected-bucket-owner](#)

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

[x-amz-request-payer](#)

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

[x-amz-server-side-encryption-customer-algorithm](#)

Specifies the algorithm to use when decrypting the object (for example, AES256).

If you encrypt an object by using server-side encryption with customer-provided encryption keys (SSE-C) when you store the object in Amazon S3, then when you GET the object, you must use the following headers:

- x-amz-server-side-encryption-customer-algorithm
- x-amz-server-side-encryption-customer-key
- x-amz-server-side-encryption-customer-key-MD5

For more information about SSE-C, see [Server-Side Encryption \(Using Customer-Provided Encryption Keys\)](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

x-amz-server-side-encryption-customer-key

Specifies the customer-provided encryption key that you originally provided for Amazon S3 to encrypt the data before storing it. This value is used to decrypt the object when recovering it and must match the one used when storing the data. The key must be appropriate for use with the algorithm specified in the `x-amz-server-side-encryption-customer-algorithm` header.

If you encrypt an object by using server-side encryption with customer-provided encryption keys (SSE-C) when you store the object in Amazon S3, then when you GET the object, you must use the following headers:

- `x-amz-server-side-encryption-customer-algorithm`
- `x-amz-server-side-encryption-customer-key`
- `x-amz-server-side-encryption-customer-key-MD5`

For more information about SSE-C, see [Server-Side Encryption \(Using Customer-Provided Encryption Keys\)](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

x-amz-server-side-encryption-customer-key-MD5

Specifies the 128-bit MD5 digest of the customer-provided encryption key according to RFC 1321. Amazon S3 uses this header for a message integrity check to ensure that the encryption key was transmitted without error.

If you encrypt an object by using server-side encryption with customer-provided encryption keys (SSE-C) when you store the object in Amazon S3, then when you GET the object, you must use the following headers:

- `x-amz-server-side-encryption-customer-algorithm`
- `x-amz-server-side-encryption-customer-key`
- `x-amz-server-side-encryption-customer-key-MD5`

For more information about SSE-C, see [Server-Side Encryption \(Using Customer-Provided Encryption Keys\)](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
x-amz-delete-marker: DeleteMarker
accept-ranges: AcceptRanges
x-amz-expiration: Expiration
x-amz-restore: Restore
Last-Modified: LastModified
Content-Length: ContentLength
ETag: ETag
x-amz-checksum-crc32: ChecksumCRC32
x-amz-checksum-crc32c: ChecksumCRC32C
x-amz-checksum-crc64nvme: ChecksumCRC64NVME
x-amz-checksum-sha1: ChecksumSHA1
x-amz-checksum-sha256: ChecksumSHA256
x-amz-checksum-sha512: ChecksumSHA512
x-amz-checksum-md5: ChecksumMD5
x-amz-checksum-xxhash64: ChecksumXXHASH64
x-amz-checksum-xxhash3: ChecksumXXHASH3
x-amz-checksum-xxhash128: ChecksumXXHASH128
x-amz-checksum-type: ChecksumType
x-amz-missing-meta: MissingMeta
x-amz-version-id: VersionId
Cache-Control: CacheControl
Content-Disposition: ContentDisposition
Content-Encoding: ContentEncoding
Content-Language: ContentLanguage
Content-Range: ContentRange
Content-Type: ContentType
Expires: Expires
x-amz-website-redirect-location: WebsiteRedirectLocation
x-amz-server-side-encryption: ServerSideEncryption
x-amz-server-side-encryption-customer-algorithm: SSECustomerAlgorithm
```

```
x-amz-server-side-encryption-customer-key-MD5: SSECustomerKeyMD5
x-amz-server-side-encryption-aws-kms-key-id: SSEKMSKeyId
x-amz-server-side-encryption-bucket-key-enabled: BucketKeyEnabled
x-amz-storage-class: StorageClass
x-amz-request-charged: RequestCharged
x-amz-replication-status: ReplicationStatus
x-amz-mp-parts-count: PartsCount
x-amz-tagging-count: TagCount
x-amz-object-lock-mode: ObjectLockMode
x-amz-object-lock-retain-until-date: ObjectLockRetainUntilDate
x-amz-object-lock-legal-hold: ObjectLockLegalHoldStatus
x-amz-object-lock-event-hold: ObjectLockEventHold
x-amz-object-lock-event-hold-duration-days: ObjectLockEventHoldDurationDays
x-amz-object-lock-event-hold-duration-years: ObjectLockEventHoldDurationYears
```

Body

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

accept-ranges

Indicates that a range of bytes was specified in the request.

Cache-Control

Specifies caching behavior along the request/reply chain.

Content-Disposition

Specifies presentational information for the object.

Content-Encoding

Indicates what content encodings have been applied to the object and thus what decoding mechanisms must be applied to obtain the media-type referenced by the Content-Type header field.

Content-Language

The language the content is in.

Content-Length

Size of the body in bytes.

Content-Range

The portion of the object returned in the response.

Content-Type

A standard MIME type describing the format of the object data.

ETag

An entity tag (ETag) is an opaque identifier assigned by a web server to a specific version of a resource found at a URL.

Expires

The date and time at which the object is no longer cacheable.

Last-Modified

Date and time when the object was last modified.

General purpose buckets - When you specify a `versionId` of the object in your request, if the specified version in the request is a delete marker, the response returns a 405 Method Not Allowed error and the `Last-Modified: timestamp` response header.

x-amz-checksum-crc32

The Base64 encoded, 32-bit CRC32 checksum of the object. This checksum is only present if the checksum was uploaded with the object. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-crc32c

The Base64 encoded, 32-bit CRC32C checksum of the object. This checksum is only present if the checksum was uploaded with the object. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-crc64nvme

The Base64 encoded, 64-bit CRC64NVME checksum of the object. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-checksum-md5

The Base64 encoded, 128-bit MD5 digest of the object. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-checksum-sha1

The Base64 encoded, 160-bit SHA1 digest of the object. This checksum is only present if the checksum was uploaded with the object. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-sha256

The Base64 encoded, 256-bit SHA256 digest of the object. This checksum is only present if the checksum was uploaded with the object. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-sha512

The Base64 encoded, 512-bit SHA512 digest of the object. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-checksum-type

The checksum type, which determines how part-level checksums are combined to create an object-level checksum for multipart objects. You can use this header response to verify that the checksum type that is received is the same checksum type that was specified in the CreateMultipartUpload request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Valid Values: COMPOSITE | FULL_OBJECT

x-amz-checksum-xxhash128

The Base64 encoded, 128-bit XXHASH128 checksum of the object. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-checksum-xxhash3

The Base64 encoded, 64-bit XXHASH3 checksum of the object. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-checksum-xxhash64

The Base64 encoded, 64-bit XXHASH64 checksum of the object. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-delete-marker

Indicates whether the object retrieved was (true) or was not (false) a Delete Marker. If false, this response header does not appear in the response.

Note

- If the current version of the object is a delete marker, Amazon S3 behaves as if the object was deleted and includes `x-amz-delete-marker: true` in the response.
- If the specified version in the request is a delete marker, the response returns a 405 Method Not Allowed error and the `Last-Modified: timestamp` response header.

x-amz-expiration

If the object expiration is configured (see [PutBucketLifecycleConfiguration](#)), the response includes this header. It includes the `expiry-date` and `rule-id` key-value pairs providing object expiration information. The value of the `rule-id` is URL-encoded.

Note

Object expiration information is not returned in directory buckets and this header returns the value "NotImplemented" in all responses for directory buckets.

x-amz-missing-meta

This is set to the number of metadata entries not returned in the headers that are prefixed with `x-amz-meta-`. This can happen if you create metadata using an API like SOAP that supports more flexible metadata than the REST API. For example, using SOAP, you can create metadata whose values are not legal HTTP headers.

Note

This functionality is not supported for directory buckets.

x-amz-mp-parts-count

The count of parts this object has. This value is only returned if you specify `partNumber` in your request and the object was uploaded as a multipart upload.

x-amz-object-lock-event-hold

The event hold status for this object. This header is only returned if the requester has the `s3:GetObjectRetention` permission.

Note

This functionality is not supported for directory buckets.

Valid Values: ON | OFF

x-amz-object-lock-event-hold-duration-days

The event hold duration in days for this object. Only returned when the event hold is enabled.

Note

This functionality is not supported for directory buckets.

x-amz-object-lock-event-hold-duration-years

The event hold duration in years for this object. Only returned when the event hold is enabled.

Note

This functionality is not supported for directory buckets.

x-amz-object-lock-legal-hold

Indicates whether this object has an active legal hold. This field is only returned if you have permission to view an object's legal hold status.

Note

This functionality is not supported for directory buckets.

Valid Values: ON | OFF

x-amz-object-lock-mode

The Object Lock mode that's currently in place for this object.

Note

This functionality is not supported for directory buckets.

Valid Values: GOVERNANCE | COMPLIANCE

x-amz-object-lock-retain-until-date

The date and time when this object's Object Lock will expire.

Note

This functionality is not supported for directory buckets.

x-amz-replication-status

Amazon S3 can return this if your request involves a bucket that is either a source or destination in a replication rule.

Note

This functionality is not supported for directory buckets.

Valid Values: COMPLETE | PENDING | FAILED | REPLICA | COMPLETED

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

x-amz-restore

Provides information about object restoration action and expiration time of the restored object copy.

Note

This functionality is not supported for directory buckets. Directory buckets only support EXPRESS_ONEZONE (the S3 Express One Zone storage class) in Availability Zones and ONEZONE_IA (the S3 One Zone-Infrequent Access storage class) in Dedicated Local Zones.

x-amz-server-side-encryption

The server-side encryption algorithm used when you store this object in Amazon S3 or Amazon FSx.

Note

When accessing data stored in Amazon FSx file systems using S3 access points, the only valid server side encryption option is `aws:fsx`.

Valid Values: AES256 | aws:fsx | aws:backup | aws:kms | aws:kms:dsse

x-amz-server-side-encryption-aws-kms-key-id

If present, indicates the ID of the KMS key that was used for object encryption.

x-amz-server-side-encryption-bucket-key-enabled

Indicates whether the object uses an S3 Bucket Key for server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS).

x-amz-server-side-encryption-customer-algorithm

If server-side encryption with a customer-provided encryption key was requested, the response will include this header to confirm the encryption algorithm that's used.

Note

This functionality is not supported for directory buckets.

x-amz-server-side-encryption-customer-key-MD5

If server-side encryption with a customer-provided encryption key was requested, the response will include this header to provide the round-trip message integrity verification of the customer-provided encryption key.

Note

This functionality is not supported for directory buckets.

x-amz-storage-class

Provides storage class information of the object. Amazon S3 returns this header for all objects except for S3 Standard storage class objects.

Note

Directory buckets - Directory buckets only support EXPRESS_ONEZONE (the S3 Express One Zone storage class) in Availability Zones and ONEZONE_IA (the S3 One Zone-Infrequent Access storage class) in Dedicated Local Zones.

Valid Values: STANDARD | REDUCED_REDUNDANCY | STANDARD_IA | ONEZONE_IA | INTELLIGENT_TIERING | GLACIER | DEEP_ARCHIVE | OUTPOSTS | GLACIER_IR |

SNOW | EXPRESS_ONEZONE | FSX_OPENZFS | FSX_ONTAP | AWS_BACKUP_WARM |
AWS_BACKUP_LOW_COST_WARM

x-amz-tagging-count

The number of tags, if any, on the object, when you have the relevant permission to read object tags.

You can use [GetObjectTagging](#) to retrieve the tag set associated with an object.

Note

This functionality is not supported for directory buckets.

x-amz-version-id

Version ID of the object.

Note

This functionality is not supported for directory buckets.

x-amz-website-redirect-location

If the bucket is configured as a website, redirects requests for this object to another object in the same bucket or to an external URL. Amazon S3 stores the value of this header in the object metadata.

Note

This functionality is not supported for directory buckets.

The following data is returned in binary format by the service.

Body

Errors

InvalidObjectState

Object is archived and inaccessible until restored.

If the object you are retrieving is stored in the S3 Glacier Flexible Retrieval storage class, the S3 Glacier Deep Archive storage class, the S3 Intelligent-Tiering Archive Access tier, or the S3 Intelligent-Tiering Deep Archive Access tier, before you can retrieve the object you must first restore a copy using [RestoreObject](#). Otherwise, this operation returns an `InvalidObjectState` error. For information about restoring archived objects, see [Restoring Archived Objects](#) in the *Amazon S3 User Guide*.

HTTP Status Code: 403

NoSuchKey

The specified key does not exist.

HTTP Status Code: 404

Examples

Sample Request for general purpose buckets

The following request returns the object `my-image.jpg`.

```
GET /my-image.jpg HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
Date: Mon, 3 Oct 2016 22:32:00 GMT
Authorization: authorization string
```

Sample Response for general purpose buckets

This example illustrates one usage of `GetObject`.

```
HTTP/1.1 200 OK
x-amz-id-2:
eftixk72aD6Ap51TnqcoF8eFidJG9Z/2mkiDFu8yU9AS1ed40pIszj7UDNEHGran
```

```
x-amz-request-id: 318BC8BC148832E5
Date: Mon, 3 Oct 2016 22:32:00 GMT
Last-Modified: Wed, 12 Oct 2009 17:50:00 GMT
ETag: "fba9dede5f27731c9771645a39863328"
Content-Length: 434234
```

[434234 bytes of object data]

Sample Response for general purpose buckets: Object with associated tags

If the object had tags associated with it, Amazon S3 returns the `x-amz-tagging-count` header with tag count.

```
HTTP/1.1 200 OK
x-amz-id-2:
eftixk72aD6Ap51TnqcoF8eFidJG9Z/2mkiDFu8yU9AS1ed40pIszj7UDNEHGran
x-amz-request-id: 318BC8BC148832E5
Date: Mon, 3 Oct 2016 22:32:00 GMT
Last-Modified: Wed, 12 Oct 2009 17:50:00 GMT
ETag: "fba9dede5f27731c9771645a39863328"
Content-Length: 434234
x-amz-tagging-count: 2
```

[434234 bytes of object data]

Sample Response for general purpose buckets: Object with an expiration

If the object had expiration set using lifecycle configuration, you get the following response with the `x-amz-expiration` header.

```
HTTP/1.1 200 OK
x-amz-id-2:
eftixk72aD6Ap51TnqcoF8eFidJG9Z/2mkiDFu8yU9AS1ed40pIszj7UDNEHGran
x-amz-request-id: 318BC8BC148832E5
Date: Wed, 28 Oct 2009 22:32:00 GMT
Last-Modified: Wed, 12 Oct 2009 17:50:00 GMT
x-amz-expiration: expiry-date="Fri, 23 Dec 2012 00:00:00 GMT", rule-
id="picture-deletion-rule"
ETag: "fba9dede5f27731c9771645a39863328"
```

```
Content-Length: 434234
Content-Type: text/plain

[434234 bytes of object data]
```

Sample Response for general purpose buckets: If an object is archived in the S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive storage classes

If the object you are retrieving is stored in the S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive storage classes, you must first restore a copy using [RestoreObject](#). Otherwise, this action returns an `InvalidObjectState` error.

```
HTTP/1.1 403 Forbidden
x-amz-request-id: CD4BD8A1310A11B3
x-amz-id-2: m9RDbQU0+RRBTj0UN1ChQ1eqMUnr9dv8b
+KP6I2gHfRJZSTSrMCoRP8RtPRzX9mb
Content-Type: application/xml
Date: Mon, 12 Nov 2012 23:53:21 GMT
Server: Amazon S3
Content-Length: 231

<Error>
  <Code>InvalidObjectState</Code>
  <Message>The action is not valid for the object's storage class</Message>
  <RequestId>9FEFFF118E15B86F</RequestId>
  <HostId>WVQ5kzhiT+oiUfDC0i0Yv8W4Tk9eNcxWi/MK+hTS/av34Xy4rBU3zsavf0aaaaa</
HostId>
</Error>
```

Sample Response for general purpose buckets: If an object is archived with the S3 Intelligent-Tiering Archive or S3 Intelligent-Tiering Deep Archive tiers

If the object you are retrieving is stored in the S3 Intelligent-Tiering Archive or S3 Intelligent-Tiering Deep Archive tiers, you must first restore a copy using [RestoreObject](#). Otherwise, this action returns an `InvalidObjectState` error. When restoring from Archive Access or Deep Archive Access tiers, the response will include `StorageClass` and `AccessTier` elements. Access tier valid values are `ARCHIVE_ACCESS` and `DEEP_ARCHIVE_ACCESS`. There is no syntax change if there is an ongoing restore.

```

HTTP/1.1 403 Forbidden
x-amz-request-id: CB6AW8C4332B23B7
x-amz-id-2: n3RRfT90+PJDUhut3nhGW2ehfhfNU5f55c
+a2ceCC36ab7c7fe3a71Q273b9Q45b1R5
Content-Type: application/xml
Date: Mon, 12 Nov 2012 23:53:21 GMT
Server: Amazon S3
Content-Length: 231

<Error>
  <Code>InvalidObjectState</Code>
  <Message>The action is not valid for the object's access tier</Message>
  <StorageClass>INTELLIGENT_TIERING</StorageClass>
  <AccessTier>ARCHIVE_ACCESS</AccessTier>
  <RequestId>9FEFFF118E15B86F</RequestId>
  <HostId>WVQ5kzhiT+oiUfDC0i0Yv8W4Tk9eNcxWi/MK+hTS/av34Xy4rBU3zsavf0aaaaa</
HostId>
</Error>

```

Sample Response for general purpose buckets: If the Latest Object Is a Delete Marker

Notice that the delete marker returns a 404 Not Found error.

```

HTTP/1.1 404 Not Found
x-amz-request-id: 318BC8BC148832E5
x-amz-id-2: eftixk72aD6Ap51Tnqzj7UDNEHGran
x-amz-version-id: 3GL4kqtJlcpXroDTDm3vjVBH40Nr8X8g
x-amz-delete-marker: true
Date: Wed, 28 Oct 2009 22:32:00 GMT
Content-Type: text/plain
Connection: close
Server: AmazonS3

```

Sample Request for general purpose buckets: Getting a specified version of an object

The following request returns the specified version of an object.

```
GET /myObject?versionId=3/L4kqtJlcpXroDTDmpUMLUo HTTP/1.1
```



```
Host: bucket.s3.<Region>.amazonaws.com
Date: Wed, 28 Oct 2009 22:32:00 GMT
Authorization: authorization string
```

Sample Response for general purpose buckets: GET a versioned object

This example illustrates one usage of `GetObject`.

```
HTTP/1.1 200 OK
x-amz-id-2: eftixk72aD6Ap540pIszj7UDNEHGran
x-amz-request-id: 318BC8BC148832E5
Date: Wed, 28 Oct 2009 22:32:00 GMT
Last-Modified: Sun, 1 Jan 2006 12:00:00 GMT
x-amz-version-id: 3/L4kqtJlcpXroDTDmJ+rmSpXd3QBpUMLUo
ETag: "fba9dede5f27731c9771645a39863328"
Content-Length: 434234
Content-Type: text/plain
Connection: close
Server: AmazonS3
[434234 bytes of object data]
```

Sample Request for general purpose buckets: Parameters altering response header values

The following request specifies all the query string parameters in a GET request overriding the response header values.

```
GET /Junk3.txt?response-cache-control=No-cache&response-content-
disposition=attachment%3B%20filename%3Dtesting.txt&response-content-encoding=x-
gzip&response-content-language=mi%2C%20en&response-expires=Thu%2C%2001%20Dec
%201994%2016:00:00%20GMT HTTP/1.1
x-amz-date: Sun, 19 Dec 2010 01:53:44 GMT
Accept: */*
Authorization: AWS AKIAIOSFODNN7EXAMPLE:aaStE6nKnw8ihhiIdReoXYlMamW=
```

Sample Response for general purpose buckets: With overridden response header values

The following request specifies all the query string parameters in a GET request overriding the response header values.

```

HTTP/1.1 200 OK
x-amz-id-2: SIidWAK3hK+Il3/
Qqiu1ZKEuegzLAAspwsqwnwygb9GgFseeFHL5CII8NXSrfWW2
x-amz-request-id: 881B1CBD9DF17WA1
Date: Sun, 19 Dec 2010 01:54:01 GMT
x-amz-meta-param1: value 1
x-amz-meta-param2: value 2
Cache-Control: No-cache
Content-Language: mi, en
Expires: Thu, 01 Dec 1994 16:00:00 GMT
Content-Disposition: attachment; filename=testing.txt
Content-Encoding: x-gzip
Last-Modified: Fri, 17 Dec 2010 18:10:41 GMT
ETag: "0332bee1a7bf845f176c5c0d1ae7cf07"
Accept-Ranges: bytes
Content-Type: text/plain
Content-Length: 22
Server: AmazonS3

[object data not shown]

```

Sample Request for general purpose buckets: Range header

The following request specifies the HTTP Range header to retrieve the first 10 bytes of an object. For more information about the HTTP Range header, see <https://www.rfc-editor.org/rfc/rfc9110.html#name-range>.

Note

Amazon S3 doesn't support retrieving multiple ranges of data per GET request.

```

GET /example-object HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
x-amz-date: Fri, 28 Jan 2011 21:32:02 GMT
Range: bytes=0-9
Authorization: AWS AKIAIOSFODNN7EXAMPLE:Yxg83MZaEgh30Z3l0rLo5RTX11o=
Sample Response with Specified Range of the Object Bytes

```

Sample Response for general purpose buckets

In the following sample response, note that the header values are set to the values specified in the true request.

```
HTTP/1.1 206 Partial Content
x-amz-id-2: MzRIS0wyjmnupCzjI1WC06l5TTAzm7/JypPGXLh00VFGcJaa03KW/
hRAqK0pIEEp
x-amz-request-id: 47622117804B3E11
Date: Fri, 28 Jan 2011 21:32:09 GMT
x-amz-meta-title: the title
Last-Modified: Fri, 28 Jan 2011 20:10:32 GMT
ETag: "b2419b1e3fd45d596ee22bdf62aaaa2f"
Accept-Ranges: bytes
Content-Range: bytes 0-9/443
Content-Type: text/plain
Content-Length: 10
Server: AmazonS3

[10 bytes of object data]
```

Sample Request for general purpose buckets: Get an object stored using server-side encryption with customer-provided encryption keys

If an object is stored in Amazon S3 using server-side encryption with customer-provided encryption keys, Amazon S3 needs encryption information so that it can decrypt the object before sending it to you in response to a GET request. You provide the encryption information in your GET request using the relevant headers, as shown in the following example request.

```
GET /example-object HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com

Accept: */*
Authorization: authorization string
Date: Wed, 28 May 2014 19:24:44 +0000
x-amz-server-side-encryption-customer-
key:g0lCfA3Dv40jZz5SQJ1ZukLRFqtI5WorC/8SEKEXAMPLE
```

```
x-amz-server-side-encryption-customer-key-MD5:ZjQrne1X/iTcskbY2m3example
x-amz-server-side-encryption-customer-algorithm:AES256
```

Sample Response for general purpose buckets

The following sample response shows some of the response headers Amazon S3 returns. Note that it includes the encryption information in the response.

```
HTTP/1.1 200 OK
x-amz-id-2: ka5jRm8X3N12ZiY29Z989zg2tNSJPMcK+to7jNjxImXBbyChqc6tLAv
+sau7Vjzh
x-amz-request-id: 195157E3E073D3F9
Date: Wed, 28 May 2014 19:24:45 GMT
Last-Modified: Wed, 28 May 2014 19:21:01 GMT
ETag: "c12022c9a3c6d3a28d29d90933a2b096"
x-amz-server-side-encryption-customer-algorithm: AES256
x-amz-server-side-encryption-customer-key-MD5: ZjQrne1X/iTcskbY2m3example
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetObjectAcl

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Returns the access control list (ACL) of an object. To use this operation, you must have `s3:GetObjectAcl` permissions or `READ_ACP` access to the object. For more information, see [Mapping of ACL permissions and access policy permissions](#) in the *Amazon S3 User Guide*

This functionality is not supported for Amazon S3 on Outposts.

By default, GET returns ACL information about the current version of an object. To return ACL information about a different version, use the `versionId` subresource.

Note

If your bucket uses the bucket owner enforced setting for S3 Object Ownership, requests to read ACLs are still supported and return the `bucket-owner-full-control` ACL with the owner being the account that created the bucket. For more information, see [Controlling object ownership and disabling ACLs](#) in the *Amazon S3 User Guide*.

The following operations are related to `GetObjectAcl`:

- [GetObject](#)
- [GetObjectAttributes](#)
- [DeleteObject](#)
- [PutObject](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /{Key+}?acl&versionId=VersionId HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-request-payer: RequestPayer
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name that contains the object for which to get the ACL information.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form *AccessPointName-AccountId.s3-accesspoint.Region.amazonaws.com*. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

Required: Yes

Key

The key of the object for which to get the ACL information.

Length Constraints: Minimum length of 1.

Required: Yes

versionId

Version ID used to reference a specific version of the object.

Note

This functionality is not supported for directory buckets.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
x-amz-request-charged: RequestCharged
<?xml version="1.0" encoding="UTF-8"?>
<AccessControlPolicy>
  <Owner>
    <DisplayName>string</DisplayName>
    <ID>string</ID>
  </Owner>
  <AccessControlList>
    <Grant>
      <Grantee>
        <DisplayName>string</DisplayName>
        <EmailAddress>string</EmailAddress>
        <ID>string</ID>
```

```
<xsi:type>string</xsi:type>
<URI>string</URI>
</Grantee>
<Permission>string</Permission>
</Grant>
</AccessControlList>
</AccessControlPolicy>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

The following data is returned in XML format by the service.

AccessControlPolicy

Root level tag for the AccessControlPolicy parameters.

Required: Yes

Grants

A list of grants.

Type: Array of [Grant](#) data types

Owner

Container for the bucket owner's ID.

Type: [Owner](#) data type

Errors

NoSuchKey

The specified key does not exist.

HTTP Status Code: 404

Examples

Sample Request

The following request returns information, including the ACL, of the object `my-image.jpg`.

```
GET /my-image.jpg?acl HTTP/1.1
Host: bucket.s3.<Region>.amazonaws.com
Date: Wed, 28 Oct 2009 22:32:00 GMT
Authorization: authorization string
```

Sample Response

This example illustrates one usage of `GetObjectAcl`.

```
HTTP/1.1 200 OK
x-amz-id-2:
eftixk72aD6Ap51TnqcoF8eFidJG9Z/2mkiDFu8yU9AS1ed40pIszj7UDNEHGran
x-amz-request-id: 318BC8BC148832E5
x-amz-version-id: 4HL4kqtJlcpXroDTDmJ+rmSpXd3dIbrHY+MTRCxf3vjVBH40Nrjfkd
Date: Wed, 28 Oct 2009 22:32:00 GMT
Last-Modified: Sun, 1 Jan 2006 12:00:00 GMT
Content-Length: 124
Content-Type: text/plain
Connection: close
Server: AmazonS3

<AccessControlPolicy>
```

```

ID>
    <Owner>
      <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeebf76c078efc7c6caea54ba06a</
    </Owner>
    <AccessControlList>
      <Grant>
        <Grantee xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
          <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeebf76c078efc7c6caea54ba06a</ID>
          <Type>CanonicalUser</Type>
        </Grantee>
        <Permission>FULL_CONTROL</Permission>
      </Grant>
    </AccessControlList>
  </AccessControlPolicy>

```

Sample Request: Getting the ACL of the specific version of an object

The following request returns information, including the ACL, of the specified version of the object, my-image.jpg.

```

HTTP/1.1 GET /my-image.jpg?versionId=3/L4kqtJlcpXroDVBH40Nr8X8gdRQBpUMLUo&acl

Host: bucket.s3.<Region>.amazonaws.com
Date: Wed, 28 Oct 2009 22:32:00 GMT
Authorization: authorization string

```

Sample Response: Showing the ACL of the specific version

This example illustrates one usage of GetObjectAcl.

```

HTTP/1.1 200 OK
x-amz-id-2:
eftixk72aD6Ap51TnqcoF8eFidJG9Z/2mkiDFu8yU9AS1ed40pIszj7UDNEHGran
x-amz-request-id: 318BC8BC148832E5
Date: Wed, 28 Oct 2009 22:32:00 GMT
Last-Modified: Sun, 1 Jan 2006 12:00:00 GMT
x-amz-version-id: 3/L4kqtJlcpXroDTDmJ+rmSpXd3dIbrHY
+MTRCxf3vjVBH40Nr8X8gdRQBpUMLUo

```

```

Content-Length: 124
Content-Type: text/plain
Connection: close
Server: AmazonS3

ID>
<AccessControlPolicy>
  <Owner>
    <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeebf76c078efc7c6caea54ba06a</
  </Owner>
  <AccessControlList>
    <Grant>
      <Grantee xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
        <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeebf76c078efc7c6caea54ba06a</ID>
        <Type>CanonicalUser</Type>
      </Grantee>
      <Permission>FULL_CONTROL</Permission>
    </Grant>
  </AccessControlList>
</AccessControlPolicy>

```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetObjectAnnotation

Service: Amazon S3

Retrieves an annotation from an Amazon S3 object. To use this operation, you must have the `s3:GetObjectAnnotation` permission.

If checksum mode is enabled via the `x-amz-checksum-mode` header, Amazon S3 returns the stored checksum in the response headers for client-side validation.

Note

Annotations are not supported by the following features: S3 Inventory Reports, API Gateway, S3 Storage Lens, Amazon S3 File Gateway, Amazon FSx, S3 on Outposts, and S3 Express One Zone (directory buckets).

The following operations are related to `GetObjectAnnotation`:

- [PutObjectAnnotation](#)
- [ListObjectAnnotations](#)
- [DeleteObjectAnnotation](#)

Request Syntax

```
GET /{Key+}?annotation&annotationName=AnnotationName&versionId=VersionId HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-request-payer: RequestPayer
x-amz-expected-bucket-owner: ExpectedBucketOwner
x-amz-checksum-mode: ChecksumMode
```

URI Request Parameters

The request uses the following URI parameters.

annotationName

The name of the annotation to retrieve.

Length Constraints: Minimum length of 1. Maximum length of 512 bytes.

Required: Yes

Bucket

The name of the bucket that contains the object.

Required: Yes

Key

The object key.

Length Constraints: Minimum length of 1.

Required: Yes

versionId

The version ID of the object.

x-amz-checksum-mode

Set to ENABLED to validate the checksum of the annotation payload on retrieval.

Valid Values: ENABLED

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the bucket is owned by a different account, the request fails with an HTTP 403 (Access Denied) error.

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
x-amz-object-version-id: ObjectVersionId
Last-Modified: LastModified
Content-Length: ContentLength
ETag: ETag
x-amz-checksum-crc32: ChecksumCRC32
x-amz-checksum-crc32c: ChecksumCRC32C
x-amz-checksum-crc64nvme: ChecksumCRC64NVME
x-amz-checksum-sha1: ChecksumSHA1
x-amz-checksum-sha256: ChecksumSHA256
x-amz-checksum-sha512: ChecksumSHA512
x-amz-checksum-md5: ChecksumMD5
x-amz-checksum-xxhash64: ChecksumXXHASH64
x-amz-checksum-xxhash3: ChecksumXXHASH3
x-amz-checksum-xxhash128: ChecksumXXHASH128
x-amz-checksum-type: ChecksumType
x-amz-server-side-encryption: ServerSideEncryption
x-amz-request-charged: RequestCharged
x-amz-replication-status: ReplicationStatus

AnnotationPayload
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

Content-Length

The size of the annotation payload, in bytes.

ETag

The entity tag of the annotation.

Last-Modified

The date and time the annotation was last modified.

[x-amz-checksum-crc32](#)

The CRC32 checksum of the annotation payload.

[x-amz-checksum-crc32c](#)

The CRC32C checksum of the annotation payload.

[x-amz-checksum-crc64nvme](#)

The CRC64NVME checksum of the annotation payload.

[x-amz-checksum-md5](#)

The MD5 checksum of the annotation payload.

[x-amz-checksum-sha1](#)

The SHA1 checksum of the annotation payload.

[x-amz-checksum-sha256](#)

The SHA256 checksum of the annotation payload.

[x-amz-checksum-sha512](#)

The SHA512 checksum of the annotation payload.

[x-amz-checksum-type](#)

The type of checksum used.

Valid Values: COMPOSITE | FULL_OBJECT

[x-amz-checksum-xxhash128](#)

The XXHASH128 checksum of the annotation payload.

[x-amz-checksum-xxhash3](#)

The XXHASH3 checksum of the annotation payload.

[x-amz-checksum-xxhash64](#)

The XXHASH64 checksum of the annotation payload.

[x-amz-object-version-id](#)

The version ID of the object that the annotation is attached to.

x-amz-replication-status

The replication status of the annotation. Possible values include PENDING, COMPLETED, FAILED, and REPLICA.

Valid Values: COMPLETE | PENDING | FAILED | REPLICA | COMPLETED

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

x-amz-server-side-encryption

The server-side encryption algorithm used.

Valid Values: AES256 | aws:fsx | aws:backup | aws:kms | aws:kms:dsse

The following data is returned in binary format by the service.

AnnotationPayload

Errors

NoSuchAnnotation

The specified annotation does not exist on this object.

HTTP Status Code: 404

NoSuchBucket

The specified bucket does not exist.

HTTP Status Code: 404

NoSuchKey

The specified key does not exist.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetObjectAttributes

Service: Amazon S3

Retrieves all of the metadata from an object without returning the object itself. This operation is useful if you're interested only in an object's metadata.

GetObjectAttributes combines the functionality of HeadObject and ListParts. All of the data returned with both of those individual calls can be returned with a single call to GetObjectAttributes.

Note

Directory buckets - For directory buckets, you must make requests for this API operation to the Zonal endpoint. These endpoints support virtual-hosted-style requests in the format `https://amzn-s3-demo-bucket.s3express-zone-id.region-code.amazonaws.com/key-name` . Path-style requests are not supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

Permissions

- **General purpose bucket permissions** - To use GetObjectAttributes, you must have READ access to the object.

The other permissions that you need to use this operation depend on whether the bucket is versioned and if a version ID is passed in the GetObjectAttributes request.

- If you pass a version ID in your request, you need both the `s3:GetObjectVersion` and `s3:GetObjectVersionAttributes` permissions.
- If you do not pass a version ID in your request, you need the `s3:GetObject` and `s3:GetObjectAttributes` permissions.

For more information, see [Specifying Permissions in a Policy](#) in the *Amazon S3 User Guide*.

If the object that you request does not exist, the error Amazon S3 returns depends on whether you also have the `s3:ListBucket` permission.

- If you have the `s3:ListBucket` permission on the bucket, Amazon S3 returns an HTTP status code `404 Not Found` ("no such key") error.
- If you don't have the `s3:ListBucket` permission, Amazon S3 returns an HTTP status code `403 Forbidden` ("access denied") error.
- **Directory bucket permissions** - To grant access to this API operation on a directory bucket, we recommend that you use the [CreateSession](#) API operation for session-based authorization. Specifically, you grant the `s3express:CreateSession` permission to the directory bucket in a bucket policy or an IAM identity-based policy. Then, you make the `CreateSession` API call on the bucket to obtain a session token. With the session token in your request header, you can make API requests to this operation. After the session token expires, you make another `CreateSession` API call to generate a new session token for use. AWS CLI or SDKs create session and refresh the session token automatically to avoid service interruptions when a session expires. For more information about authorization, see [CreateSession](#).

If the object is encrypted with SSE-KMS, you must also have the `kms:GenerateDataKey` and `kms:Decrypt` permissions in IAM identity-based policies and AWS KMS key policies for the AWS KMS key.

Encryption

Note

Encryption request headers, like `x-amz-server-side-encryption`, should not be sent for HEAD requests if your object uses server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS), dual-layer server-side encryption with AWS KMS keys (DSSE-KMS), or server-side encryption with Amazon S3 managed encryption keys (SSE-S3). The `x-amz-server-side-encryption` header is used when you PUT an object to S3 and want to specify the encryption method. If you include this header in a GET request for an object that uses these types of keys, you'll get an HTTP `400 Bad Request` error. It's because the encryption method can't be changed when you retrieve the object.

If you encrypted an object when you stored the object in Amazon S3 by using server-side encryption with customer-provided encryption keys (SSE-C), then when you retrieve the

metadata from the object, you must use the following headers. These headers provide the server with the encryption key required to retrieve the object's metadata. The headers are:

- `x-amz-server-side-encryption-customer-algorithm`
- `x-amz-server-side-encryption-customer-key`
- `x-amz-server-side-encryption-customer-key-MD5`

For more information about SSE-C, see [Server-Side Encryption \(Using Customer-Provided Encryption Keys\)](#) in the *Amazon S3 User Guide*.

 Note

Directory bucket permissions - For directory buckets, there are only two supported options for server-side encryption: server-side encryption with Amazon S3 managed keys (SSE-S3) (AES256) and server-side encryption with AWS KMS keys (SSE-KMS) (`aws:kms`). We recommend that the bucket's default encryption uses the desired encryption configuration and you don't override the bucket default encryption in your `CreateSession` requests or `PUT` object requests. Then, new objects are automatically encrypted with the desired encryption settings. For more information, see [Protecting data with server-side encryption](#) in the *Amazon S3 User Guide*. For more information about the encryption overriding behaviors in directory buckets, see [Specifying server-side encryption with AWS KMS for new object uploads](#).

Versioning

Directory buckets - S3 Versioning isn't enabled and supported for directory buckets. For this API operation, only the `null` value of the version ID is supported by directory buckets. You can only specify `null` to the `versionId` query parameter in the request.

Conditional request headers

Consider the following when using request headers:

- If both of the `If-Match` and `If-Unmodified-Since` headers are present in the request as follows, then Amazon S3 returns the HTTP status code `200 OK` and the data requested:
 - `If-Match` condition evaluates to `true`.
 - `If-Unmodified-Since` condition evaluates to `false`.

For more information about conditional requests, see [RFC 7232](#).

- If both of the `If-None-Match` and `If-Modified-Since` headers are present in the request as follows, then Amazon S3 returns the HTTP status code 304 Not Modified:
 - `If-None-Match` condition evaluates to false.
 - `If-Modified-Since` condition evaluates to true.

For more information about conditional requests, see [RFC 7232](#).

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `Bucket-name.s3express-zone-id.region-code.amazonaws.com`.

The following actions are related to `GetObjectAttributes`:

- [GetObject](#)
- [GetObjectAcl](#)
- [GetObjectLegalHold](#)
- [GetObjectLockConfiguration](#)
- [GetObjectRetention](#)
- [GetObjectTagging](#)
- [HeadObject](#)
- [ListParts](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /{Key+}?attributes&versionId=VersionId HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-max-parts: MaxParts
x-amz-part-number-marker: PartNumberMarker
x-amz-server-side-encryption-customer-algorithm: SSECustomerAlgorithm
```

```
x-amz-server-side-encryption-customer-key: SSECustomerKey  
x-amz-server-side-encryption-customer-key-MD5: SSECustomerKeyMD5  
x-amz-request-payer: RequestPayer  
x-amz-expected-bucket-owner: ExpectedBucketOwner  
x-amz-object-attributes: ObjectAttributes
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket that contains the object.

Directory buckets - When you use this operation with a directory bucket, you must use virtual-hosted-style requests in the format `Bucket-name.s3express-zone-id.region-code.amazonaws.com`. Path-style requests are not supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must follow the format `bucket-base-name--zone-id--x-s3` (for example, `amzn-s3-demo-bucket--usw2-az1--x-s3`). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form `AccessPointName-AccountId.s3-accesspoint.Region.amazonaws.com`. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

Note

Object Lambda access points are not supported by directory buckets.

S3 on Outposts - When you use this action with S3 on Outposts, you must direct requests to the S3 on Outposts hostname. The S3 on Outposts hostname takes the form `AccessPointName-AccountId.outpostID.s3-outposts.Region.amazonaws.com`.

When you use this action with S3 on Outposts, the destination bucket must be the Outposts access point ARN or the access point alias. For more information about S3 on Outposts, see [What is S3 on Outposts?](#) in the *Amazon S3 User Guide*.

Required: Yes

Key

The object key.

Length Constraints: Minimum length of 1.

Required: Yes

versionId

The version ID used to reference a specific version of the object.

Note

S3 Versioning isn't enabled and supported for directory buckets. For this API operation, only the null value of the version ID is supported by directory buckets. You can only specify null to the versionId query parameter in the request.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-max-parts

Sets the maximum number of parts to return. For more information, see [Uploading and copying objects using multipart upload in Amazon S3](#) in the *Amazon Simple Storage Service user guide*.

x-amz-object-attributes

Specifies the fields at the root level that you want returned in the response. Fields that you do not specify are not returned.

Valid Values: ETag | Checksum | ObjectParts | StorageClass | ObjectSize

Required: Yes

x-amz-part-number-marker

Specifies the part after which listing should begin. Only parts with higher part numbers will be listed. For more information, see [Uploading and copying objects using multipart upload in Amazon S3](#) in the *Amazon Simple Storage Service user guide*.

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

x-amz-server-side-encryption-customer-algorithm

Specifies the algorithm to use when encrypting the object (for example, AES256).

Note

This functionality is not supported for directory buckets.

x-amz-server-side-encryption-customer-key

Specifies the customer-provided encryption key for Amazon S3 to use in encrypting data. This value is used to store the object and then it is discarded; Amazon S3 does not store the encryption key. The key must be appropriate for use with the algorithm specified in the x-amz-server-side-encryption-customer-algorithm header.

Note

This functionality is not supported for directory buckets.

x-amz-server-side-encryption-customer-key-MD5

Specifies the 128-bit MD5 digest of the encryption key according to RFC 1321. Amazon S3 uses this header for a message integrity check to ensure that the encryption key was transmitted without error.

Note

This functionality is not supported for directory buckets.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
x-amz-delete-marker: DeleteMarker
Last-Modified: LastModified
x-amz-version-id: VersionId
x-amz-request-charged: RequestCharged
<?xml version="1.0" encoding="UTF-8"?>
<GetObjectAttributesResponse>
  <ETag>string</ETag>
  <Checksum>
    <ChecksumCRC32>string</ChecksumCRC32>
    <ChecksumCRC32C>string</ChecksumCRC32C>
    <ChecksumCRC64NVME>string</ChecksumCRC64NVME>
    <ChecksumMD5>string</ChecksumMD5>
    <ChecksumSHA1>string</ChecksumSHA1>
    <ChecksumSHA256>string</ChecksumSHA256>
    <ChecksumSHA512>string</ChecksumSHA512>
    <ChecksumType>string</ChecksumType>
    <ChecksumXXHASH128>string</ChecksumXXHASH128>
    <ChecksumXXHASH3>string</ChecksumXXHASH3>
    <ChecksumXXHASH64>string</ChecksumXXHASH64>
  </Checksum>
  <ObjectParts>
    <IsTruncated>boolean</IsTruncated>
    <MaxParts>integer</MaxParts>
    <NextPartNumberMarker>integer</NextPartNumberMarker>
```

```

<PartNumberMarker>integer</PartNumberMarker>
<Part>
  <ChecksumCRC32>string</ChecksumCRC32>
  <ChecksumCRC32C>string</ChecksumCRC32C>
  <ChecksumCRC64NVME>string</ChecksumCRC64NVME>
  <ChecksumMD5>string</ChecksumMD5>
  <ChecksumSHA1>string</ChecksumSHA1>
  <ChecksumSHA256>string</ChecksumSHA256>
  <ChecksumSHA512>string</ChecksumSHA512>
  <ChecksumXXHASH128>string</ChecksumXXHASH128>
  <ChecksumXXHASH3>string</ChecksumXXHASH3>
  <ChecksumXXHASH64>string</ChecksumXXHASH64>
  <PartNumber>integer</PartNumber>
  <Size>long</Size>
</Part>
...
<PartsCount>integer</PartsCount>
</ObjectParts>
<StorageClass>string</StorageClass>
<ObjectSize>long</ObjectSize>
</GetObjectAttributesResponse>

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

Last-Modified

Date and time when the object was last modified.

x-amz-delete-marker

Specifies whether the object retrieved was (`true`) or was not (`false`) a delete marker. If `false`, this response header does not appear in the response. To learn more about delete markers, see [Working with delete markers](#).

Note

This functionality is not supported for directory buckets.

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

x-amz-version-id

The version ID of the object.

Note

This functionality is not supported for directory buckets.

The following data is returned in XML format by the service.

GetObjectAttributesResponse

Root level tag for the GetObjectAttributesResponse parameters.

Required: Yes

Checksum

The checksum or digest of the object.

Type: [Checksum](#) data type

ETag

An ETag is an opaque identifier assigned by a web server to a specific version of a resource found at a URL.

Type: String

ObjectParts

A collection of parts associated with a multipart upload.

Type: [GetObjectAttributesParts](#) data type

ObjectSize

The size of the object in bytes.

Type: Long

StorageClass

Provides the storage class information of the object. Amazon S3 returns this header for all objects except for S3 Standard storage class objects.

For more information, see [Storage Classes](#).

Note

Directory buckets - Directory buckets only support EXPRESS_ONEZONE (the S3 Express One Zone storage class) in Availability Zones and ONEZONE_IA (the S3 One Zone-Infrequent Access storage class) in Dedicated Local Zones.

Type: String

Valid Values: STANDARD | REDUCED_REDUNDANCY | STANDARD_IA | ONEZONE_IA | INTELLIGENT_TIERING | GLACIER | DEEP_ARCHIVE | OUTPOSTS | GLACIER_IR | SNOW | EXPRESS_ONEZONE | FSX_OPENZFS | FSX_ONTAP | AWS_BACKUP_WARM | AWS_BACKUP_LOW_COST_WARM

Errors

NoSuchKey

The specified key does not exist.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetObjectLegalHold

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Gets an object's current legal hold status. For more information, see [Locking Objects](#).

This functionality is not supported for Amazon S3 on Outposts.

The following action is related to GetObjectLegalHold:

- [GetObjectAttributes](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /{Key+}?legal-hold&versionId=VersionId HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-request-payer: RequestPayer
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name containing the object whose legal hold status you want to retrieve.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access

point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form *AccessPointName-AccountId.s3-accesspoint.Region.amazonaws.com*. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

Required: Yes

Key

The key name for the object whose legal hold status you want to retrieve.

Length Constraints: Minimum length of 1.

Required: Yes

versionId

The version ID of the object whose legal hold status you want to retrieve.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<LegalHold>
  <Status>string</Status>
</LegalHold>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

LegalHold

Root level tag for the LegalHold parameters.

Required: Yes

Status

Indicates whether the specified object has a legal hold in place.

Type: String

Valid Values: ON | OFF

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetObjectLockConfiguration

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Gets the Object Lock configuration for a bucket. The rule specified in the Object Lock configuration will be applied by default to every new object placed in the specified bucket. For more information, see [Locking Objects](#).

The following action is related to GetObjectLockConfiguration:

- [GetObjectAttributes](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?object-lock HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket whose Object Lock configuration you want to retrieve.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access

point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form *AccessPointName-AccountId.s3-accesspoint.Region.amazonaws.com*. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

Required: Yes

[x-amz-expected-bucket-owner](#)

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<ObjectLockConfiguration>
  <ObjectLockEnabled>string</ObjectLockEnabled>
  <Rule>
    <DefaultRetention>
      <Days>integer</Days>
      <DefaultEventHold>
        <Days>integer</Days>
        <Years>integer</Years>
      </DefaultEventHold>
      <Mode>string</Mode>
      <Years>integer</Years>
    </DefaultRetention>
  </Rule>
</ObjectLockConfiguration>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

ObjectLockConfiguration

Root level tag for the ObjectLockConfiguration parameters.

Required: Yes

ObjectLockEnabled

Indicates whether this bucket has an Object Lock configuration enabled. Enable ObjectLockEnabled when you apply ObjectLockConfiguration to a bucket.

Type: String

Valid Values: Enabled

Rule

Specifies the Object Lock rule for the specified object. Enable the this rule when you apply ObjectLockConfiguration to a bucket. Bucket settings require both a mode and a period. The period can be either Days or Years but you must select one. You cannot specify Days and Years at the same time.

Type: [ObjectLockRule](#) data type

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)

- [AWS SDK for Ruby V3](#)

GetObjectRetention

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Retrieves an object's retention settings. For more information, see [Locking Objects](#).

This functionality is not supported for Amazon S3 on Outposts.

The following action is related to GetObjectRetention:

- [GetObjectAttributes](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /{Key+}?retention&versionId=VersionId HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-request-payer: RequestPayer
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name containing the object whose retention settings you want to retrieve.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access

point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form *AccessPointName-AccountId.s3-accesspoint.Region.amazonaws.com*. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

Required: Yes

Key

The key name for the object whose retention settings you want to retrieve.

Length Constraints: Minimum length of 1.

Required: Yes

versionId

The version ID for the object whose retention settings you want to retrieve.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<Retention>
  <Mode>string</Mode>
  <RetainUntilDate>timestamp</RetainUntilDate>
  <EventHold>string</EventHold>
  <EventHoldDuration>
    <Days>integer</Days>
    <Years>integer</Years>
  </EventHoldDuration>
</Retention>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

Retention

Root level tag for the Retention parameters.

Required: Yes

EventHold

The event hold status for the object. Set to ON to enable an event hold or OFF to disable it.

Type: String

Valid Values: ON | OFF

EventHoldDuration

The event hold duration for the object. Specifies how long the object remains protected after the event hold is released.

Type: EventHoldDuration data type

Mode

Indicates the Retention mode for the specified object.

Type: String

Valid Values: GOVERNANCE | COMPLIANCE

RetainUntilDate

The date on which this Object Lock Retention will expire.

Type: Timestamp

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetObjectTagging

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Returns the tag-set of an object. You send the GET request against the tagging subresource associated with the object.

To use this operation, you must have permission to perform the `s3:GetObjectTagging` action. By default, the GET action returns information about current version of an object. For a versioned bucket, you can have multiple versions of an object in your bucket. To retrieve tags of any other version, use the `versionId` query parameter. You also need permission for the `s3:GetObjectVersionTagging` action.

By default, the bucket owner has this permission and can grant this permission to others.

For information about the Amazon S3 object tagging feature, see [Object Tagging](#).

The following actions are related to `GetObjectTagging`:

- [DeleteObjectTagging](#)
- [GetObjectAttributes](#)
- [PutObjectTagging](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /{Key+}?tagging&versionId=VersionId HTTP/1.1
Host: Bucket.s3.amazonaws.com
```

```
x-amz-expected-bucket-owner: ExpectedBucketOwner  
x-amz-request-payer: RequestPayer
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name containing the object for which to get the tagging information.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form *AccessPointName-AccountId.s3-accesspoint.Region.amazonaws.com*. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

S3 on Outposts - When you use this action with S3 on Outposts, you must direct requests to the S3 on Outposts hostname. The S3 on Outposts hostname takes the form *AccessPointName-AccountId.outpostID.s3-outposts.Region.amazonaws.com*. When you use this action with S3 on Outposts, the destination bucket must be the Outposts access point ARN or the access point alias. For more information about S3 on Outposts, see [What is S3 on Outposts?](#) in the *Amazon S3 User Guide*.

Required: Yes

Key

Object key for which to get the tagging information.

Length Constraints: Minimum length of 1.

Required: Yes

versionId

The versionId of the object for which to get the tagging information.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
x-amz-version-id: VersionId
<?xml version="1.0" encoding="UTF-8"?>
<Tagging>
  <TagSet>
    <Tag>
      <Key>string</Key>
      <Value>string</Value>
    </Tag>
  </TagSet>
</Tagging>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

x-amz-version-id

The versionId of the object for which you got the tagging information.

The following data is returned in XML format by the service.

Tagging

Root level tag for the Tagging parameters.

Required: Yes

TagSet

Contains the tag set.

Type: Array of [Tag](#) data types

Examples

Sample Request

The following request returns the tag set of the specified object.

```
GET /example-object?tagging HTTP/1.1
Host: examplebucket.s3.<Region>.amazonaws.com
Date: Thu, 22 Sep 2016 21:33:08 GMT
Authorization: authorization string
```

Sample Response

This example illustrates one usage of `GetObjectTagging`.

```
HTTP/1.1 200 OK
Date: Thu, 22 Sep 2016 21:33:08 GMT
Connection: close
Server: AmazonS3
<?xml version="1.0" encoding="UTF-8"?>
<Tagging xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <TagSet>
    <Tag>
      <Key>tag1</Key>
      <Value>val1</Value>
    </Tag>
    <Tag>
      <Key>tag2</Key>
      <Value>val2</Value>
    </Tag>
  </TagSet>
</Tagging>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetObjectTorrent

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Returns torrent files from a bucket. BitTorrent can save you bandwidth when you're distributing large files.

Note

You can get torrent only for objects that are less than 5 GB in size, and that are not encrypted using server-side encryption with a customer-provided encryption key.

To use GET, you must have READ access to the object.

This functionality is not supported for Amazon S3 on Outposts.

The following action is related to GetObjectTorrent:

- [GetObject](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /{Key+}?torrent HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-request-payer: RequestPayer
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket containing the object for which to get the torrent files.

Required: Yes

Key

The object key for which to get the information.

Length Constraints: Minimum length of 1.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
x-amz-request-charged: RequestCharged
```

Body

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

The following data is returned in binary format by the service.

Body

Examples

Getting torrent files in a bucket

This example retrieves the Torrent file for the Nelson object in the quotes bucket.

```
GET /quotes/Nelson?torrent HTTP/1.0
Host: bucket.s3.<Region>.amazonaws.com
Date: Wed, 28 Oct 2009 22:32:00 GMT
```

```
Authorization: authorization string
```

Sample Response

This example illustrates one usage of `GetObjectTorrent`.

```
HTTP/1.1 200 OK
x-amz-request-id: 7CD745EBB7AB5ED9
Date: Wed, 25 Nov 2009 12:00:00 GMT
Content-Disposition: attachment; filename=Nelson.torrent;
Content-Type: application/x-bittorrent
Content-Length: 537
Server: AmazonS3

<body: a Bencoded dictionary as defined by the BitTorrent specification>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetPublicAccessBlock

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Retrieves the `PublicAccessBlock` configuration for an Amazon S3 bucket. This operation returns the bucket-level configuration only. To understand the effective public access behavior, you must also consider account-level settings (which may inherit from organization-level policies). To use this operation, you must have the `s3:GetBucketPublicAccessBlock` permission. For more information about Amazon S3 permissions, see [Specifying Permissions in a Policy](#).

Important

When Amazon S3 evaluates the `PublicAccessBlock` configuration for a bucket or an object, it checks the `PublicAccessBlock` configuration for both the bucket (or the bucket that contains the object) and the bucket owner's account. Account-level settings automatically inherit from organization-level policies when present. If the `PublicAccessBlock` settings are different between the bucket and the account, Amazon S3 uses the most restrictive combination of the bucket-level and account-level settings.

For more information about when Amazon S3 considers a bucket or an object public, see [The Meaning of "Public"](#).

The following operations are related to `GetPublicAccessBlock`:

- [Using Amazon S3 Block Public Access](#)
- [PutPublicAccessBlock](#)
- [GetPublicAccessBlock](#)
- [DeletePublicAccessBlock](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?publicAccessBlock HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the Amazon S3 bucket whose `PublicAccessBlock` configuration you want to retrieve.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<PublicAccessBlockConfiguration>
  <BlockPublicAcls>boolean</BlockPublicAcls>
```

```
<IgnorePublicAcls>boolean</IgnorePublicAcls>  
<BlockPublicPolicy>boolean</BlockPublicPolicy>  
<RestrictPublicBuckets>boolean</RestrictPublicBuckets>  
</PublicAccessBlockConfiguration>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

PublicAccessBlockConfiguration

Root level tag for the PublicAccessBlockConfiguration parameters.

Required: Yes

BlockPublicAcls

Specifies whether Amazon S3 should block public access control lists (ACLs) for this bucket and objects in this bucket. Setting this element to TRUE causes the following behavior:

- PUT Bucket ACL and PUT Object ACL calls fail if the specified ACL is public.
- PUT Object calls fail if the request includes a public ACL.
- PUT Bucket calls fail if the request includes a public ACL.

Enabling this setting doesn't affect existing policies or ACLs.

Type: Boolean

BlockPublicPolicy

Specifies whether Amazon S3 should block public bucket policies for this bucket. Setting this element to TRUE causes Amazon S3 to reject calls to PUT Bucket policy if the specified bucket policy allows public access.

Enabling this setting doesn't affect existing bucket policies.

Type: Boolean

IgnorePublicAcls

Specifies whether Amazon S3 should ignore public ACLs for this bucket and objects in this bucket. Setting this element to TRUE causes Amazon S3 to ignore all public ACLs on this bucket and objects in this bucket.

Enabling this setting doesn't affect the persistence of any existing ACLs and doesn't prevent new public ACLs from being set.

Type: Boolean

RestrictPublicBuckets

Specifies whether Amazon S3 should restrict public bucket policies for this bucket. Setting this element to TRUE restricts access to this bucket to only AWS service principals and authorized users within this account if the bucket has a public policy.

Enabling this setting doesn't affect previously stored bucket policies, except that public and cross-account access within any public bucket policy, including non-public delegation to specific accounts, is blocked.

Type: Boolean

Examples

Sample Request

The following request gets a bucket `PublicAccessBlock` configuration.

```
GET /<bucket-name>?publicAccessBlock HTTP/1.1
Host: <bucket-name>.s3.<Region>.amazonaws.com
x-amz-date: <Thu, 15 Nov 2016 00:17:21 GMT>
Authorization: <signatureValue>
```

Sample Response

This example illustrates one usage of `GetPublicAccessBlock`.

```
HTTP/1.1 200 OK
x-amz-id-2: ITnGT1y4REXAMPLEPi4hk1TXouTf0hccUjo0iCPEXAMPLEutBj3M7fPG1W02SEWp
x-amz-request-id: 51991EXAMPLE5321
Date: Thu, 15 Nov 2016 00:17:22 GMT
Server: AmazonS3
Content-Length: 0
```

```
<PublicAccessBlockConfiguration>
  <BlockPublicAcls>TRUE</BlockPublicAcls>
  <IgnorePublicAcls>FALSE</IgnorePublicAcls>
  <BlockPublicPolicy>FALSE</BlockPublicPolicy>
  <RestrictPublicBuckets>FALSE</RestrictPublicBuckets>
</PublicAccessBlockConfiguration>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

HeadBucket

Service: Amazon S3

You can use this operation to determine if a bucket exists and if you have permission to access it. The action returns a 200 OK HTTP status code if the bucket exists and you have permission to access it. You can make a HeadBucket call on any bucket name to any Region in the partition, and regardless of the permissions on the bucket, you will receive a response header with the correct bucket location so that you can then make a proper, signed request to the appropriate Regional endpoint.

Note

If the bucket doesn't exist or you don't have permission to access it, the HEAD request returns a generic 400 Bad Request, 403 Forbidden, or 404 Not Found HTTP status code. A message body isn't included, so you can't determine the exception beyond these HTTP response codes.

Authentication and authorization

General purpose buckets - Request to public buckets that grant the s3:ListBucket permission publicly do not need to be signed. All other HeadBucket requests must be authenticated and signed by using IAM credentials (access key ID and secret access key for the IAM identities). All headers with the x-amz- prefix, including x-amz-copy-source, must be signed. For more information, see [REST Authentication](#).

Directory buckets - You must use IAM credentials to authenticate and authorize your access to the HeadBucket API operation, instead of using the temporary security credentials through the CreateSession API operation.

AWS CLI or SDKs handles authentication and authorization on your behalf.

Permissions

- **General purpose bucket permissions** - To use this operation, you must have permissions to perform the s3:ListBucket action. The bucket owner has this permission by default and can grant this permission to others. For more information about permissions, see [Managing access permissions to your Amazon S3 resources](#) in the *Amazon S3 User Guide*.

- **Directory bucket permissions** - You must have the **s3express:CreateSession** permission in the Action element of a policy. If no session mode is specified, the session will be created with the maximum allowable privilege, attempting ReadWrite first, then ReadOnly if ReadWrite is not permitted. If you want to explicitly restrict the access to be read-only, you can set the s3express:SessionMode condition key to ReadOnly on the bucket.

For more information about example bucket policies, see [Example bucket policies for S3 Express One Zone](#) and [AWS Identity and Access Management \(IAM\) identity-based policies for S3 Express One Zone](#) in the *Amazon S3 User Guide*.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `Bucket-name.s3express-zone-id.region-code.amazonaws.com`.

Note

You must make requests for this API operation to the Zonal endpoint. These endpoints support virtual-hosted-style requests in the format `https://bucket-name.s3express-zone-id.region-code.amazonaws.com`. Path-style requests are not supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after my, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
HEAD / HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name.

Directory buckets - When you use this operation with a directory bucket, you must use virtual-hosted-style requests in the format `Bucket-name.s3express-zone-id.region-code.amazonaws.com`. Path-style requests are not supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must follow the format `bucket-base-name--zone-id--x-s3` (for example, `amzn-s3-demo-bucket--usw2-az1--x-s3`). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form `AccessPointName-AccountId.s3-accesspoint.Region.amazonaws.com`. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

Object Lambda access points - When you use this API operation with an Object Lambda access point, provide the alias of the Object Lambda access point in place of the bucket name. If the Object Lambda access point alias in a request is not valid, the error code `InvalidAccessPointAliasError` is returned. For more information about `InvalidAccessPointAliasError`, see [List of Error Codes](#).

Note

Object Lambda access points are not supported by directory buckets.

S3 on Outposts - When you use this action with S3 on Outposts, you must direct requests to the S3 on Outposts hostname. The S3 on Outposts hostname takes the form

`AccessPointName-AccountId.outpostID.s3-outposts.Region.amazonaws.com`.
When you use this action with S3 on Outposts, the destination bucket must be the Outposts access point ARN or the access point alias. For more information about S3 on Outposts, see [What is S3 on Outposts?](#) in the *Amazon S3 User Guide*.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
x-amz-bucket-arn: BucketArn
x-amz-bucket-location-type: BucketLocationType
x-amz-bucket-location-name: BucketLocationName
x-amz-bucket-region: BucketRegion
x-amz-access-point-alias: AccessPointAlias
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

x-amz-access-point-alias

Indicates whether the bucket name used in the request is an access point alias.

Note

For directory buckets, the value of this field is false.

x-amz-bucket-arn

The Amazon Resource Name (ARN) of the S3 bucket. ARNs uniquely identify AWS resources across all of AWS.

Note

This parameter is only supported for S3 directory buckets. For more information, see [Using tags with directory buckets](#).

Length Constraints: Minimum length of 1. Maximum length of 128.

Pattern: `arn:[^:]+:(s3|s3express):.*`

x-amz-bucket-location-name

The name of the location where the bucket will be created.

For directory buckets, the Zone ID of the Availability Zone or the Local Zone where the bucket is created. An example Zone ID value for an Availability Zone is `usw2-az1`.

Note

This functionality is only supported by directory buckets.

x-amz-bucket-location-type

The type of location where the bucket is created.

Note

This functionality is only supported by directory buckets.

Valid Values: `AvailabilityZone` | `LocalZone`

x-amz-bucket-region

The Region that the bucket is located.

Length Constraints: Minimum length of 0. Maximum length of 20.

Errors

NoSuchBucket

The specified bucket does not exist.

HTTP Status Code: 404

Examples

Sample Request for general purpose buckets

This example illustrates one usage of HeadBucket.

```
HEAD / HTTP/1.1
Date: Fri, 10 Feb 2012 21:34:55 GMT
Authorization: authorization string
Host: myawsbucket.s3.amazonaws.com
Connection: Keep-Alive
```

Sample responses for general purpose buckets

This example illustrates one usage of HeadBucket.

```
HTTP/1.1 200 OK
x-amz-id-2: wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY
x-amz-request-id: 28PG2CEB32F5EE25
x-amz-bucket-region: us-west-2
x-amz-access-point-alias: false
Date: Fri, 10 2012 21:34:56 GMT
Server: AmazonS3
```

```
HTTP/1.1 301 Moved Permanently
```

```
x-amz-id-2: wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY
x-amz-request-id: M8H3G2AJ0V6NQ8AJ
Date: Tue, 06 May 2025 21:29:41 GMT
x-amz-bucket-region: us-east-1
x-amz-access-point-alias: false
Server: AmazonS3
```

```
HTTP/1.1 403 Forbidden
x-amz-bucket-region: us-east-1
x-amz-request-id: P6K9QSW8EN5YEPS
x-amz-id-2: wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY
Date: Tue, 06 May 2025 21:24:51 GMT
Server: AmazonS3
```

Example of anonymous request response

This example illustrates one usage of HeadBucket.

```
HTTP/1.1 403 Forbidden
x-amz-bucket-region: us-east-1
x-amz-request-id: 7B3WF90PGN6J8F2
x-amz-id-2: wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY
Date: Tue, 06 May 2025 21:55:02
GMT Server: AmazonS3
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)

- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

HeadObject

Service: Amazon S3

The HEAD operation retrieves metadata from an object without returning the object itself. This operation is useful if you're interested only in an object's metadata.

Note

A HEAD request has the same options as a GET operation on an object. The response is identical to the GET response except that there is no response body. Because of this, if the HEAD request generates an error, it returns a generic code, such as 400 Bad Request, 403 Forbidden, 404 Not Found, 405 Method Not Allowed, 412 Precondition Failed, or 304 Not Modified. It's not possible to retrieve the exact exception of these error codes.

Request headers are limited to 8 KB in size. For more information, see [Common Request Headers](#).

Permissions

- **General purpose bucket permissions** - To use HEAD, you must have the `s3:GetObject` permission. You need the relevant read object (or version) permission for this operation. For more information, see [Actions, resources, and condition keys for Amazon S3](#) in the *Amazon S3 User Guide*. For more information about the permissions to S3 API operations by S3 resource types, see [Required permissions for Amazon S3 API operations](#) in the *Amazon S3 User Guide*.

If the object you request doesn't exist, the error that Amazon S3 returns depends on whether you also have the `s3:ListBucket` permission.

- If you have the `s3:ListBucket` permission on the bucket, Amazon S3 returns an HTTP status code 404 Not Found error.
- If you don't have the `s3:ListBucket` permission, Amazon S3 returns an HTTP status code 403 Forbidden error.
- **Directory bucket permissions** - To grant access to this API operation on a directory bucket, we recommend that you use the [CreateSession](#) API operation for session-based authorization. Specifically, you grant the `s3express:CreateSession` permission to the directory bucket in a bucket policy or an IAM identity-based policy. Then, you make the `CreateSession` API call on the bucket to obtain a session token. With the session token in

your request header, you can make API requests to this operation. After the session token expires, you make another `CreateSession` API call to generate a new session token for use. AWS CLI or SDKs create session and refresh the session token automatically to avoid service interruptions when a session expires. For more information about authorization, see [CreateSession](#).

If you enable `x-amz-checksum-mode` in the request and the object is encrypted with AWS Key Management Service (AWS KMS), you must also have the `kms:GenerateDataKey` and `kms:Decrypt` permissions in IAM identity-based policies and AWS KMS key policies for the AWS KMS key to retrieve the checksum of the object.

Encryption

Note

Encryption request headers, like `x-amz-server-side-encryption`, should not be sent for HEAD requests if your object uses server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS), dual-layer server-side encryption with AWS KMS keys (DSSE-KMS), or server-side encryption with Amazon S3 managed encryption keys (SSE-S3). The `x-amz-server-side-encryption` header is used when you PUT an object to S3 and want to specify the encryption method. If you include this header in a HEAD request for an object that uses these types of keys, you'll get an HTTP 400 Bad Request error. It's because the encryption method can't be changed when you retrieve the object.

If you encrypt an object by using server-side encryption with customer-provided encryption keys (SSE-C) when you store the object in Amazon S3, then when you retrieve the metadata from the object, you must use the following headers to provide the encryption key for the server to be able to retrieve the object's metadata. The headers are:

- `x-amz-server-side-encryption-customer-algorithm`
- `x-amz-server-side-encryption-customer-key`
- `x-amz-server-side-encryption-customer-key-MD5`

For more information about SSE-C, see [Server-Side Encryption \(Using Customer-Provided Encryption Keys\)](#) in the *Amazon S3 User Guide*.

Note

Directory bucket - For directory buckets, there are only two supported options for server-side encryption: SSE-S3 and SSE-KMS. SSE-C isn't supported. For more information, see [Protecting data with server-side encryption](#) in the *Amazon S3 User Guide*.

Versioning

- If the current version of the object is a delete marker, Amazon S3 behaves as if the object was deleted and includes `x-amz-delete-marker: true` in the response.
- If the specified version is a delete marker, the response returns a 405 Method Not Allowed error and the `Last-Modified: timestamp` response header.

Note

- **Directory buckets** - Delete marker is not supported for directory buckets.
- **Directory buckets** - S3 Versioning isn't enabled and supported for directory buckets. For this API operation, only the `null` value of the version ID is supported by directory buckets. You can only specify `null` to the `versionId` query parameter in the request.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `Bucket-name.s3express-zone-id.region-code.amazonaws.com`.

Note

For directory buckets, you must make requests for this API operation to the Zonal endpoint. These endpoints support virtual-hosted-style requests in the format `https://amzn-s3-demo-bucket.s3express-zone-id.region-code.amazonaws.com/key-name`. Path-style requests are not supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more

information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

The following actions are related to HeadObject:

- [GetObject](#)
- [GetObjectAttributes](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
HEAD /Key+?partNumber=PartNumber&response-cache-control=ResponseCacheControl&response-  
content-disposition=ResponseContentDisposition&response-  
content-encoding=ResponseContentEncoding&response-content-  
language=ResponseContentLanguage&response-content-type=ResponseContentType&response-  
expires=ResponseExpires&versionId=VersionId HTTP/1.1  
Host: Bucket.s3.amazonaws.com  
If-Match: IfMatch  
If-Modified-Since: IfModifiedSince  
If-None-Match: IfNoneMatch  
If-Unmodified-Since: IfUnmodifiedSince  
Range: Range  
x-amz-server-side-encryption-customer-algorithm: SSECustomerAlgorithm  
x-amz-server-side-encryption-customer-key: SSECustomerKey  
x-amz-server-side-encryption-customer-key-MD5: SSECustomerKeyMD5  
x-amz-request-payer: RequestPayer  
x-amz-expected-bucket-owner: ExpectedBucketOwner  
x-amz-checksum-mode: ChecksumMode
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket that contains the object.

Directory buckets - When you use this operation with a directory bucket, you must use virtual-hosted-style requests in the format `Bucket-name.s3express-zone-id.region-code.amazonaws.com`. Path-style requests are not supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must follow the format `bucket-base-name--zone-id--x-s3` (for example, `amzn-s3-demo-bucket--usw2-az1--x-s3`). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form `AccessPointName-AccountId.s3-accesspoint.Region.amazonaws.com`. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

Note

Object Lambda access points are not supported by directory buckets.

S3 on Outposts - When you use this action with S3 on Outposts, you must direct requests to the S3 on Outposts hostname. The S3 on Outposts hostname takes the form `AccessPointName-AccountId.outpostID.s3-outposts.Region.amazonaws.com`. When you use this action with S3 on Outposts, the destination bucket must be the Outposts access point ARN or the access point alias. For more information about S3 on Outposts, see [What is S3 on Outposts?](#) in the *Amazon S3 User Guide*.

Required: Yes

If-Match

Return the object only if its entity tag (ETag) is the same as the one specified; otherwise, return a 412 (precondition failed) error.

If both of the `If-Match` and `If-Unmodified-Since` headers are present in the request as follows:

- `If-Match` condition evaluates to `true`, and;
- `If-Unmodified-Since` condition evaluates to `false`;

Then Amazon S3 returns `200 OK` and the data requested.

For more information about conditional requests, see [RFC 7232](#).

If-Modified-Since

Return the object only if it has been modified since the specified time; otherwise, return a `304 (not modified)` error.

If both of the `If-None-Match` and `If-Modified-Since` headers are present in the request as follows:

- `If-None-Match` condition evaluates to `false`, and;
- `If-Modified-Since` condition evaluates to `true`;

Then Amazon S3 returns the `304 Not Modified` response code.

For more information about conditional requests, see [RFC 7232](#).

If-None-Match

Return the object only if its entity tag (ETag) is different from the one specified; otherwise, return a `304 (not modified)` error.

If both of the `If-None-Match` and `If-Modified-Since` headers are present in the request as follows:

- `If-None-Match` condition evaluates to `false`, and;
- `If-Modified-Since` condition evaluates to `true`;

Then Amazon S3 returns the `304 Not Modified` response code.

For more information about conditional requests, see [RFC 7232](#).

If-Unmodified-Since

Return the object only if it has not been modified since the specified time; otherwise, return a `412 (precondition failed)` error.

If both of the If-Match and If-Unmodified-Since headers are present in the request as follows:

- If-Match condition evaluates to true, and;
- If-Unmodified-Since condition evaluates to false;

Then Amazon S3 returns 200 OK and the data requested.

For more information about conditional requests, see [RFC 7232](#).

Key

The object key.

Length Constraints: Minimum length of 1.

Required: Yes

partNumber

Part number of the object being read. This is a positive integer between 1 and 10,000.

Effectively performs a 'ranged' HEAD request for the part specified. Useful querying about the size of the part and the number of parts in this object.

Range

HeadObject returns only the metadata for an object. If the Range is satisfiable, only the ContentLength is affected in the response. If the Range is not satisfiable, S3 returns a 416 - Requested Range Not Satisfiable error.

response-cache-control

Sets the Cache-Control header of the response.

response-content-disposition

Sets the Content-Disposition header of the response.

response-content-encoding

Sets the Content-Encoding header of the response.

response-content-language

Sets the Content-Language header of the response.

response-content-type

Sets the Content-Type header of the response.

response-expires

Sets the Expires header of the response.

versionId

Version ID used to reference a specific version of the object.

Note

For directory buckets in this API operation, only the null value of the version ID is supported.

x-amz-checksum-mode

To retrieve the checksum, this parameter must be enabled.

General purpose buckets - If you enable checksum mode and the object is uploaded with a [checksum](#) and encrypted with an AWS Key Management Service (AWS KMS) key, you must have permission to use the `kms:Decrypt` action to retrieve the checksum.

Directory buckets - If you enable `ChecksumMode` and the object is encrypted with AWS Key Management Service (AWS KMS), you must also have the `kms:GenerateDataKey` and `kms:Decrypt` permissions in IAM identity-based policies and AWS KMS key policies for the AWS KMS key to retrieve the checksum of the object.

Valid Values: ENABLED

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For

information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: `requester`

[x-amz-server-side-encryption-customer-algorithm](#)

Specifies the algorithm to use when encrypting the object (for example, AES256).

Note

This functionality is not supported for directory buckets.

[x-amz-server-side-encryption-customer-key](#)

Specifies the customer-provided encryption key for Amazon S3 to use in encrypting data. This value is used to store the object and then it is discarded; Amazon S3 does not store the encryption key. The key must be appropriate for use with the algorithm specified in the `x-amz-server-side-encryption-customer-algorithm` header.

Note

This functionality is not supported for directory buckets.

[x-amz-server-side-encryption-customer-key-MD5](#)

Specifies the 128-bit MD5 digest of the encryption key according to RFC 1321. Amazon S3 uses this header for a message integrity check to ensure that the encryption key was transmitted without error.

Note

This functionality is not supported for directory buckets.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
x-amz-delete-marker: DeleteMarker
accept-ranges: AcceptRanges
x-amz-expiration: Expiration
x-amz-restore: Restore
x-amz-archive-status: ArchiveStatus
Last-Modified: LastModified
Content-Length: ContentLength
x-amz-checksum-crc32: ChecksumCRC32
x-amz-checksum-crc32c: ChecksumCRC32C
x-amz-checksum-crc64nvme: ChecksumCRC64NVME
x-amz-checksum-sha1: ChecksumSHA1
x-amz-checksum-sha256: ChecksumSHA256
x-amz-checksum-sha512: ChecksumSHA512
x-amz-checksum-md5: ChecksumMD5
x-amz-checksum-xxhash64: ChecksumXXHASH64
x-amz-checksum-xxhash3: ChecksumXXHASH3
x-amz-checksum-xxhash128: ChecksumXXHASH128
x-amz-checksum-type: ChecksumType
ETag: ETag
x-amz-missing-meta: MissingMeta
x-amz-version-id: VersionId
Cache-Control: CacheControl
Content-Disposition: ContentDisposition
Content-Encoding: ContentEncoding
Content-Language: ContentLanguage
Content-Type: ContentType
Content-Range: ContentRange
Expires: Expires
x-amz-website-redirect-location: WebsiteRedirectLocation
x-amz-server-side-encryption: ServerSideEncryption
x-amz-server-side-encryption-customer-algorithm: SSECustomerAlgorithm
x-amz-server-side-encryption-customer-key-MD5: SSECustomerKeyMD5
x-amz-server-side-encryption-aws-kms-key-id: SSEKMSKeyId
x-amz-server-side-encryption-bucket-key-enabled: BucketKeyEnabled
x-amz-storage-class: StorageClass
x-amz-request-charged: RequestCharged
```

```
x-amz-replication-status: ReplicationStatus  
x-amz-mp-parts-count: PartsCount  
x-amz-tagging-count: TagCount  
x-amz-object-lock-mode: ObjectLockMode  
x-amz-object-lock-retain-until-date: ObjectLockRetainUntilDate  
x-amz-object-lock-legal-hold: ObjectLockLegalHoldStatus  
x-amz-object-lock-event-hold: ObjectLockEventHold  
x-amz-object-lock-event-hold-duration-days: ObjectLockEventHoldDurationDays  
x-amz-object-lock-event-hold-duration-years: ObjectLockEventHoldDurationYears
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

accept-ranges

Indicates that a range of bytes was specified.

Cache-Control

Specifies caching behavior along the request/reply chain.

Content-Disposition

Specifies presentational information for the object.

Content-Encoding

Indicates what content encodings have been applied to the object and thus what decoding mechanisms must be applied to obtain the media-type referenced by the Content-Type header field.

Content-Language

The language the content is in.

Content-Length

Size of the body in bytes.

Content-Range

The portion of the object returned in the response for a GET request.

Content-Type

A standard MIME type describing the format of the object data.

ETag

An entity tag (ETag) is an opaque identifier assigned by a web server to a specific version of a resource found at a URL.

Expires

The date and time at which the object is no longer cacheable.

Last-Modified

Date and time when the object was last modified.

x-amz-archive-status

The archive state of the head object.

Note

This functionality is not supported for directory buckets.

Valid Values: `ARCHIVE_ACCESS` | `DEEP_ARCHIVE_ACCESS`

x-amz-checksum-crc32

The Base64 encoded, 32-bit CRC32 checksum of the object. This checksum is only present if the checksum was uploaded with the object. When you use an API operation on an object that was uploaded using multipart uploads, this value may not be a direct checksum value of the full object. Instead, it's a calculation based on the checksum values of each individual part. For more information about how checksums are calculated with multipart uploads, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-crc32c

The Base64 encoded, 32-bit CRC32C checksum of the object. This checksum is only present if the checksum was uploaded with the object. When you use an API operation on an object that was uploaded using multipart uploads, this value may not be a direct checksum value of the full object. Instead, it's a calculation based on the checksum values of each individual part. For more

information about how checksums are calculated with multipart uploads, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-crc64nvme

The Base64 encoded, 64-bit CRC64NVME checksum of the object. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-checksum-md5

The Base64 encoded, 128-bit MD5 digest of the object. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-checksum-sha1

The Base64 encoded, 160-bit SHA1 digest of the object. This checksum is only present if the checksum was uploaded with the object. When you use the API operation on an object that was uploaded using multipart uploads, this value may not be a direct checksum value of the full object. Instead, it's a calculation based on the checksum values of each individual part. For more information about how checksums are calculated with multipart uploads, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-sha256

The Base64 encoded, 256-bit SHA256 digest of the object. This checksum is only present if the checksum was uploaded with the object. When you use an API operation on an object that was uploaded using multipart uploads, this value may not be a direct checksum value of the full object. Instead, it's a calculation based on the checksum values of each individual part. For more information about how checksums are calculated with multipart uploads, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-sha512

The Base64 encoded, 512-bit SHA512 digest of the object. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-checksum-type

The checksum type, which determines how part-level checksums are combined to create an object-level checksum for multipart objects. You can use this header response to verify that the checksum type that is received is the same checksum type that was specified in CreateMultipartUpload request. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

Valid Values: COMPOSITE | FULL_OBJECT

x-amz-checksum-xxhash128

The Base64 encoded, 128-bit XXHASH128 checksum of the object. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-checksum-xxhash3

The Base64 encoded, 64-bit XXHASH3 checksum of the object. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-checksum-xxhash64

The Base64 encoded, 64-bit XXHASH64 checksum of the object. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-delete-marker

Specifies whether the object retrieved was (true) or was not (false) a Delete Marker. If false, this response header does not appear in the response.

Note

This functionality is not supported for directory buckets.

x-amz-expiration

If the object expiration is configured (see [PutBucketLifecycleConfiguration](#)), the response includes this header. It includes the expiry-date and rule-id key-value pairs providing object expiration information. The value of the rule-id is URL-encoded.

Note

Object expiration information is not returned in directory buckets and this header returns the value "NotImplemented" in all responses for directory buckets.

x-amz-missing-meta

This is set to the number of metadata entries not returned in x-amz-meta headers. This can happen if you create metadata using an API like SOAP that supports more flexible metadata

than the REST API. For example, using SOAP, you can create metadata whose values are not legal HTTP headers.

Note

This functionality is not supported for directory buckets.

x-amz-mp-parts-count

The count of parts this object has. This value is only returned if you specify `partNumber` in your request and the object was uploaded as a multipart upload.

x-amz-object-lock-event-hold

The event hold status for this object. This header is only returned if the requester has the `s3:GetObjectRetention` permission.

Note

This functionality is not supported for directory buckets.

Valid Values: ON | OFF

x-amz-object-lock-event-hold-duration-days

The event hold duration in days for this object. Only returned when the event hold is enabled.

Note

This functionality is not supported for directory buckets.

x-amz-object-lock-event-hold-duration-years

The event hold duration in years for this object. Only returned when the event hold is enabled.

Note

This functionality is not supported for directory buckets.

x-amz-object-lock-legal-hold

Specifies whether a legal hold is in effect for this object. This header is only returned if the requester has the `s3:GetObjectLegalHold` permission. This header is not returned if the specified version of this object has never had a legal hold applied. For more information about S3 Object Lock, see [Object Lock](#).

Note

This functionality is not supported for directory buckets.

Valid Values: ON | OFF

x-amz-object-lock-mode

The Object Lock mode, if any, that's in effect for this object. This header is only returned if the requester has the `s3:GetObjectRetention` permission. For more information about S3 Object Lock, see [Object Lock](#).

Note

This functionality is not supported for directory buckets.

Valid Values: GOVERNANCE | COMPLIANCE

x-amz-object-lock-retain-until-date

The date and time when the Object Lock retention period expires. This header is only returned if the requester has the `s3:GetObjectRetention` permission.

Note

This functionality is not supported for directory buckets.

x-amz-replication-status

Amazon S3 can return this header if your request involves a bucket that is either a source or a destination in a replication rule.

In replication, you have a source bucket on which you configure replication and destination bucket or buckets where Amazon S3 stores object replicas. When you request an object (GetObject) or object metadata (HeadObject) from these buckets, Amazon S3 will return the `x-amz-replication-status` header in the response as follows:

- **If requesting an object from the source bucket**, Amazon S3 will return the `x-amz-replication-status` header if the object in your request is eligible for replication.

For example, suppose that in your replication configuration, you specify object prefix `TaxDocs` requesting Amazon S3 to replicate objects with key prefix `TaxDocs`. Any objects you upload with this key name prefix, for example `TaxDocs/document1.pdf`, are eligible for replication. For any object request with this key name prefix, Amazon S3 will return the `x-amz-replication-status` header with value `PENDING`, `COMPLETED` or `FAILED` indicating object replication status.

- **If requesting an object from a destination bucket**, Amazon S3 will return the `x-amz-replication-status` header with value `REPLICA` if the object in your request is a replica that Amazon S3 created and there is no replica modification replication in progress.
- **When replicating objects to multiple destination buckets**, the `x-amz-replication-status` header acts differently. The header of the source object will only return a value of `COMPLETED` when replication is successful to all destinations. The header will remain at value `PENDING` until replication has completed for all destinations. If one or more destinations fails replication the header will return `FAILED`.

For more information, see [Replication](#).

 Note

This functionality is not supported for directory buckets.

Valid Values: `COMPLETE` | `PENDING` | `FAILED` | `REPLICA` | `COMPLETED`

[x-amz-request-charged](#)

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

x-amz-restore

If the object is an archived object (an object whose storage class is GLACIER), the response includes this header if either the archive restoration is in progress (see [RestoreObject](#) or an archive copy is already restored.

If an archive copy is already restored, the header value indicates when Amazon S3 is scheduled to delete the object copy. For example:

```
x-amz-restore: ongoing-request="false", expiry-date="Fri, 21 Dec 2012
00:00:00 GMT"
```

If the object restoration is in progress, the header returns the value ongoing-request="true".

For more information about archiving objects, see [Transitioning Objects: General Considerations](#).

Note

This functionality is not supported for directory buckets. Directory buckets only support EXPRESS_ONEZONE (the S3 Express One Zone storage class) in Availability Zones and ONEZONE_IA (the S3 One Zone-Infrequent Access storage class) in Dedicated Local Zones.

x-amz-server-side-encryption

The server-side encryption algorithm used when you store this object in Amazon S3 or Amazon FSx.

Note

When accessing data stored in Amazon FSx file systems using S3 access points, the only valid server side encryption option is `aws:fsx`.

Valid Values: AES256 | `aws:fsx` | `aws:backup` | `aws:kms` | `aws:kms:dsse`

x-amz-server-side-encryption-aws-kms-key-id

If present, indicates the ID of the KMS key that was used for object encryption.

x-amz-server-side-encryption-bucket-key-enabled

Indicates whether the object uses an S3 Bucket Key for server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS).

x-amz-server-side-encryption-customer-algorithm

If server-side encryption with a customer-provided encryption key was requested, the response will include this header to confirm the encryption algorithm that's used.

Note

This functionality is not supported for directory buckets.

x-amz-server-side-encryption-customer-key-MD5

If server-side encryption with a customer-provided encryption key was requested, the response will include this header to provide the round-trip message integrity verification of the customer-provided encryption key.

Note

This functionality is not supported for directory buckets.

x-amz-storage-class

Provides storage class information of the object. Amazon S3 returns this header for all objects except for S3 Standard storage class objects.

For more information, see [Storage Classes](#).

Note

Directory buckets - Directory buckets only support EXPRESS_ONEZONE (the S3 Express One Zone storage class) in Availability Zones and ONEZONE_IA (the S3 One Zone-Infrequent Access storage class) in Dedicated Local Zones.

Valid Values: STANDARD | REDUCED_REDUNDANCY | STANDARD_IA | ONEZONE_IA | INTELLIGENT_TIERING | GLACIER | DEEP_ARCHIVE | OUTPOSTS | GLACIER_IR | SNOW | EXPRESS_ONEZONE | FSX_OPENZFS | FSX_ONTAP | AWS_BACKUP_WARM | AWS_BACKUP_LOW_COST_WARM

x-amz-tagging-count

The number of tags, if any, on the object, when you have the relevant permission to read object tags.

You can use [GetObjectTagging](#) to retrieve the tag set associated with an object.

Note

This functionality is not supported for directory buckets.

x-amz-version-id

Version ID of the object.

Note

This functionality is not supported for directory buckets.

x-amz-website-redirect-location

If the bucket is configured as a website, redirects requests for this object to another object in the same bucket or to an external URL. Amazon S3 stores the value of this header in the object metadata.

Note

This functionality is not supported for directory buckets.

Errors

NoSuchKey

The specified key does not exist.

HTTP Status Code: 404

Examples

Sample Request for general purpose buckets

The following request returns the metadata of an object.

```
HEAD /my-image.jpg HTTP/1.1
Host: bucket.s3.<Region>.amazonaws.com
Date: Wed, 28 Oct 2009 22:32:00 GMT
Authorization: AWS AKIAIOSFODNN7EXAMPLE:02236Q3V0RonhpaBX5sCYVf1bNRuU=
```

Sample Response for general purpose buckets

This example illustrates one usage of HeadObject.

```
HTTP/1.1 200 OK
x-amz-id-2: ef8yU9AS1ed40pIszj7UDNEHGran
x-amz-request-id: 318BC8BC143432E5
x-amz-version-id: 3HL4kqtJlcpXroDTDmjVBH40NrjfkD
Date: Wed, 28 Oct 2009 22:32:00 GMT
Last-Modified: Sun, 1 Jan 2006 12:00:00 GMT
ETag: "fba9dede5f27731c9771645a39863328"
Content-Length: 434234
Content-Type: text/plain
Connection: close
```

```
Server: AmazonS3
```

Sample Response for general purpose buckets: With an expiration tag

If the object is scheduled to expire according to a lifecycle configuration set on the bucket, the response returns the `x-amz-expiration` tag with information about when Amazon S3 will delete the object. For more information, see [Transitioning Objects: General Considerations](#).

```
HTTP/1.1 200 OK
x-amz-id-2: azQRZtQJ2m1P8R+TIsG9h0VuC/DmiSJmjXUMq7snk
+LKSJeurtmfzSlGhR46GzSJ
x-amz-request-id: 0EFF61CCE3F24A26
Date: Mon, 17 Dec 2012 02:26:39 GMT
Last-Modified: Mon, 17 Dec 2012 02:14:10 GMT
x-amz-expiration: expiry-date="Fri, 21 Dec 2012 00:00:00 GMT", rule-
id="Rule for testfile.txt"
ETag: "54b0c58c7ce9f2a8b551351102ee0938"
Accept-Ranges: bytes
Content-Type: text/plain
Content-Length: 14
Server: AmazonS3
```

Sample Request for general purpose buckets: Getting metadata from a specified version of an object

The following request returns the metadata of the specified version of an object.

```
HEAD /my-image.jpg?versionId=3HL4kqCxf3vjVBH40Nrjfk HTTP/1.1
Host: bucket.s3.<Region>.amazonaws.com
Date: Wed, 28 Oct 2009 22:32:00 GMT
Authorization: AWS AKIAIOSFODNN7EXAMPLE:02236Q3V0WpaBX5sCYVf1bNRuU=
```

Sample Response for general purpose buckets: To a versioned HEAD request

This example illustrates one usage of `HeadObject`.

```

HTTP/1.1 200 OK
x-amz-id-2: eftixk72aD6Ap51TnqcoF8epIszj7UDNEHGran
x-amz-request-id: 318BC8BC143432E5
x-amz-version-id: 3HL4kqtJlcpXrof3vjVBH40Nrjfkd
Date: Wed, 28 Oct 2009 22:32:00 GMT
Last-Modified: Sun, 1 Jan 2006 12:00:00 GMT
ETag: "fba9dede5f27731c9771645a39863328"
Content-Length: 434234
Content-Type: text/plain
Connection: close
Server: AmazonS3

```

Sample Request for general purpose buckets: For an S3 Glacier object

For an archived object, the `x-amz-restore` header provides the date when the restored copy expires, as shown in the following response. Even if the object is stored in S3 Glacier, all object metadata is still available.

```

HEAD /my-image.jpg HTTP/1.1
Host: bucket.s3.<Region>.amazonaws.com
Date: 13 Nov 2012 00:28:38 GMT
Authorization: AWS AKIAIOSFODNN7EXAMPLE:02236Q3V0RonhpaBX5sCYVf1bNRuU=

```

Sample Response for general purpose buckets: S3 Glacier object

If the object is already restored, the `x-amz-restore` header provides the date when the restored copy will expire, as shown in the following response.

```

HTTP/1.1 200 OK
x-amz-id-2: FSVaTMjrmBp3Izs1NnwBZeu7M19iI8UbxMbi0A8AirHANJBo
+hEftBuiESACOMJp
x-amz-request-id: E5CEFCB143EB505A
Date: Tue, 13 Nov 2012 00:28:38 GMT
Last-Modified: Mon, 15 Oct 2012 21:58:07 GMT
x-amz-restore: ongoing-request="false", expiry-date="Wed, 07 Nov 2012
00:00:00 GMT"
ETag: "1accb31fcf202eba0c0f41fa2f09b4d7"
Accept-Ranges: bytes

```

```
Content-Type: binary/octet-stream
Content-Length: 300
Server: AmazonS3
```

Sample Response for general purpose buckets: In-progress restoration

If the restoration is in progress, the `x-amz-restore` header returns a message accordingly.

```
HTTP/1.1 200 OK
x-amz-id-2: b+V2mDiMHTdy1myoUBpctvmJl95H9U/OSUm/
jRtHxjh0+pCk5SvByL4xu2TDv4GM
x-amz-request-id: E2E7B6AEE4E9BD2B
Date: Tue, 13 Nov 2012 00:43:32 GMT
Last-Modified: Sat, 20 Oct 2012 21:28:27 GMT
x-amz-restore: ongoing-request="true"
ETag: "1accb31fcf202eba0c0f41fa2f09b4d7"
Accept-Ranges: bytes
Content-Type: binary/octet-stream
Content-Length: 300
Server: AmazonS3
```

Sample Response for general purpose buckets: Object archived using S3 Intelligent-Tiering

If an object is stored using the S3 Intelligent-Tiering storage class and is currently in one of the archive tiers, then this action shows the current tier using the `x-amz-archive-status` header.

```
HTTP/1.1 200 OK
x-amz-id-2: FSVaTMjrmBp3Izs1NnwBZeu7M19iI8UbxMbi0A8AirHANJBo
+hEftBuiESACOMJp
x-amz-request-id: E5CEFCB143EB505A
Date: Fri, 13 Nov 2020 00:28:38 GMT
Last-Modified: Mon, 15 Oct 2012 21:58:07 GMT
ETag: "1accb31fcf202eba0c0f41fa2f09b4d7"
x-amz-storage-class: 'INTELLIGENT_TIERING'
x-amz-archive-status: 'ARCHIVE_ACCESS'
Accept-Ranges: bytes
Content-Type: binary/octet-stream
Content-Length: 300
Server: AmazonS3
```

Sample Response for general purpose buckets: Object archived using S3 Intelligent-Tiering with restore in progress

If an object is stored using the S3 Intelligent-Tiering storage class and is currently in the process of being restored from one of the archive tiers, then this action shows the current tier using the `x-amz-archive-status` header and the current restore status using the `x-amz-restore` header.

```
HTTP/1.1 200 OK
x-amz-id-2: FSVaTMjrmBp3Izs1NnwBZeu7M19iI8UbxMbi0A8AirHANJBo
+hEftBuiESACOMJp
x-amz-request-id: E5CEFCB143EB505A
Date: Fri, 13 Nov 2020 00:28:38 GMT
Last-Modified: Mon, 15 Oct 2012 21:58:07 GMT
ETag: "1accb31fcf202eba0c0f41fa2f09b4d7"
x-amz-storage-class: 'INTELLIGENT_TIERING'
x-amz-archive-status: 'ARCHIVE_ACCESS'
x-amz-restore: 'ongoing-request="true"'
x-amz-restore-request-date: 'Fri, 13 Nov 2020 00:20:00 GMT'
Accept-Ranges: bytes
Content-Type: binary/octet-stream
Content-Length: 300
Server: AmazonS3
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)

- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListBucketAnalyticsConfigurations

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Lists the analytics configurations for the bucket. You can have up to 1,000 analytics configurations per bucket.

This action supports list pagination and does not return more than 100 configurations at a time. You should always check the `IsTruncated` element in the response. If there are no more configurations to list, `IsTruncated` is set to false. If there are more configurations to list, `IsTruncated` is set to true, and there will be a value in `NextContinuationToken`. You use the `NextContinuationToken` value to continue the pagination of the list by passing the value in `continuation-token` in the request to GET the next page.

To use this operation, you must have permissions to perform the `s3:GetAnalyticsConfiguration` action. The bucket owner has this permission by default. The bucket owner can grant this permission to others. For more information about permissions, see [Permissions Related to Bucket Subresource Operations](#) and [Managing Access Permissions to Your Amazon S3 Resources](#).

For information about Amazon S3 analytics feature, see [Amazon S3 Analytics – Storage Class Analysis](#).

The following operations are related to `ListBucketAnalyticsConfigurations`:

- [GetBucketAnalyticsConfiguration](#)
- [DeleteBucketAnalyticsConfiguration](#)
- [PutBucketAnalyticsConfiguration](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?analytics&continuation-token=ContinuationToken HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket from which analytics configurations are retrieved.

Required: Yes

continuation-token

The ContinuationToken that represents a placeholder from where this request should begin.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

```

<?xml version="1.0" encoding="UTF-8"?>
<ListBucketAnalyticsConfigurationResult>
  <IsTruncated>boolean</IsTruncated>
  <ContinuationToken>string</ContinuationToken>
  <NextContinuationToken>string</NextContinuationToken>
  <AnalyticsConfiguration>
    <Filter>
      <And>
        <Prefix>string</Prefix>
        <Tag>
          <Key>string</Key>
          <Value>string</Value>
        </Tag>
        ...
      </And>
      <Prefix>string</Prefix>
      <Tag>
        <Key>string</Key>
        <Value>string</Value>
      </Tag>
    </Filter>
    <Id>string</Id>
    <StorageClassAnalysis>
      <DataExport>
        <Destination>
          <S3BucketDestination>
            <Bucket>string</Bucket>
            <BucketAccountId>string</BucketAccountId>
            <Format>string</Format>
            <Prefix>string</Prefix>
          </S3BucketDestination>
        </Destination>
        <OutputSchemaVersion>string</OutputSchemaVersion>
      </DataExport>
    </StorageClassAnalysis>
  </AnalyticsConfiguration>
  ...
</ListBucketAnalyticsConfigurationResult>

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

ListBucketAnalyticsConfigurationResult

Root level tag for the ListBucketAnalyticsConfigurationResult parameters.

Required: Yes

AnalyticsConfiguration

The list of analytics configurations for a bucket.

Type: Array of [AnalyticsConfiguration](#) data types

ContinuationToken

The marker that is used as a starting point for this analytics configuration list response. This value is present if it was sent in the request.

Type: String

IsTruncated

Indicates whether the returned list of analytics configurations is complete. A value of true indicates that the list is not complete and the NextContinuationToken will be provided for a subsequent request.

Type: Boolean

NextContinuationToken

NextContinuationToken is sent when isTruncated is true, which indicates that there are more analytics configurations to list. The next request must include this NextContinuationToken. The token is obfuscated and is not a usable value.

Type: String

Examples

Sample Request

Delete the metric configuration with a specified ID, which disables the CloudWatch metrics with the ExampleMetrics value for the FilterId dimension.

```
GET /?analytics HTTP/1.1
```

```
Host: example-bucket.s3.<Region>.amazonaws.com
x-amz-date: 20160430T233541Z
Authorization: authorization string
```

Sample Response

This example illustrates one usage of ListBucketAnalyticsConfigurations.

```
HTTP/1.1 200 OK
x-amz-id-2: gyB+3jRPnrkN98ZajxHXr3u7EFM67bNgSAXexeEHndCX/7GRnfTXxReKUQF28IfP
x-amz-request-id: 3B3C7C725673C630
Date: Sat, 30 Apr 2016 23:29:37 GMT
Content-Length: length
Server: AmazonS3

<ListBucketAnalyticsConfigurationResult xmlns="http://s3.amazonaws.com/
doc/2006-03-01/">
  <AnalyticsConfiguration>
    <Id>list1</Id>
    <Filter>
      <And>
        <Prefix>images/</Prefix>
        <Tag>
          <Key>dog</Key>
          <Value>corgi</Value>
        </Tag>
      </And>
    </Filter>
    <StorageClassAnalysis>
      <DataExport>
        <OutputSchemaVersion>V_1</OutputSchemaVersion>
        <Destination>
          <S3BucketDestination>
            <Format>CSV</Format>
            <BucketAccountId>123456789012</BucketAccountId>
            <Bucket>arn:aws:s3:::destination-bucket</Bucket>
            <Prefix>destination-prefix</Prefix>
          </S3BucketDestination>
        </Destination>
      </DataExport>
    </StorageClassAnalysis>
  </AnalyticsConfiguration>
```

```

<AnalyticsConfiguration>
  <Id>report1</Id>
  <Filter>
    <And>
      <Prefix>images/</Prefix>
      <Tag>
        <Key>dog</Key>
        <Value>bulldog</Value>
      </Tag>
    </And>
  </Filter>
  <StorageClassAnalysis>
    <DataExport>
      <OutputSchemaVersion>V_1</OutputSchemaVersion>
      <Destination>
        <S3BucketDestination>
          <Format>CSV</Format>
          <BucketAccountId>123456789012</BucketAccountId>
          <Bucket>arn:aws:s3:::destination-bucket</Bucket>
          <Prefix>destination-prefix</Prefix>
        </S3BucketDestination>
      </Destination>
    </DataExport>
  </StorageClassAnalysis>
</AnalyticsConfiguration>
...
<IsTruncated>>false</IsTruncated>
<!-- If ContinuationToken was provided in the request. -->
<ContinuationToken>...</ContinuationToken>
<!-- if IsTruncated == true -->
<IsTruncated>>true</IsTruncated>
<NextContinuationToken>...</NextContinuationToken>
</ListBucketAnalyticsConfigurationResult>

```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)

- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListBucketIntelligentTieringConfigurations

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Lists the S3 Intelligent-Tiering configuration from the specified bucket.

The S3 Intelligent-Tiering storage class is designed to optimize storage costs by automatically moving data to the most cost-effective storage access tier, without performance impact or operational overhead. S3 Intelligent-Tiering delivers automatic cost savings in three low latency and high throughput access tiers. To get the lowest storage cost on data that can be accessed in minutes to hours, you can choose to activate additional archiving capabilities.

The S3 Intelligent-Tiering storage class is the ideal storage class for data with unknown, changing, or unpredictable access patterns, independent of object size or retention period. If the size of an object is less than 128 KB, it is not monitored and not eligible for auto-tiering. Smaller objects can be stored, but they are always charged at the Frequent Access tier rates in the S3 Intelligent-Tiering storage class.

For more information, see [Storage class for automatically optimizing frequently and infrequently accessed objects](#).

Operations related to ListBucketIntelligentTieringConfigurations include:

- [DeleteBucketIntelligentTieringConfiguration](#)
- [PutBucketIntelligentTieringConfiguration](#)
- [GetBucketIntelligentTieringConfiguration](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?intelligent-tiering&continuation-token=ContinuationToken HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the Amazon S3 bucket whose configuration you want to modify or retrieve.

Required: Yes

continuation-token

The ContinuationToken that represents a placeholder from where this request should begin.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<ListBucketIntelligentTieringConfigurationsOutput>
  <IsTruncated>boolean</IsTruncated>
  <ContinuationToken>string</ContinuationToken>
  <NextContinuationToken>string</NextContinuationToken>
  <IntelligentTieringConfiguration>
    <Filter>
      <And>
        <Prefix>string</Prefix>
```

```

    <Tag>
      <Key>string</Key>
      <Value>string</Value>
    </Tag>
    ...
  </And>
  <Prefix>string</Prefix>
  <Tag>
    <Key>string</Key>
    <Value>string</Value>
  </Tag>
</Filter>
<Id>string</Id>
<Status>string</Status>
<Tiering>
  <AccessTier>string</AccessTier>
  <Days>integer</Days>
</Tiering>
...
</IntelligentTieringConfiguration>
...
</ListBucketIntelligentTieringConfigurationsOutput>

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

ListBucketIntelligentTieringConfigurationsOutput

Root level tag for the ListBucketIntelligentTieringConfigurationsOutput parameters.

Required: Yes

ContinuationToken

The ContinuationToken that represents a placeholder from where this request should begin.

Type: String

IntelligentTieringConfiguration

The list of S3 Intelligent-Tiering configurations for a bucket.

Type: Array of [IntelligentTieringConfiguration](#) data types

IsTruncated

Indicates whether the returned list of analytics configurations is complete. A value of `true` indicates that the list is not complete and the `NextContinuationToken` will be provided for a subsequent request.

Type: Boolean

NextContinuationToken

The marker used to continue this inventory configuration listing. Use the `NextContinuationToken` from this response to continue the listing in a subsequent request. The continuation token is an opaque value that Amazon S3 understands.

Type: String

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListBucketInventoryConfigurations

Service: Amazon S3

Returns a list of S3 Inventory configurations for the bucket. You can have up to 1,000 inventory configurations per bucket.

This action supports list pagination and does not return more than 100 configurations at a time. Always check the `IsTruncated` element in the response. If there are no more configurations to list, `IsTruncated` is set to false. If there are more configurations to list, `IsTruncated` is set to true, and there is a value in `NextContinuationToken`. You use the `NextContinuationToken` value to continue the pagination of the list by passing the value in `continuation-token` in the request to GET the next page.

Note

Directory buckets - For directory buckets, you must make requests for this API operation to the Regional endpoint. These endpoints support path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name`. Virtual-hosted-style requests aren't supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

Permissions

To use this operation, you must have permissions to perform the `s3:GetInventoryConfiguration` action. The bucket owner has this permission by default. The bucket owner can grant this permission to others. For more information about permissions, see [Permissions Related to Bucket Subresource Operations](#) and [Managing Access Permissions to Your Amazon S3 Resources](#).

- **General purpose bucket permissions** - The `s3:GetInventoryConfiguration` permission is required in a policy. For more information about general purpose buckets permissions, see [Using Bucket Policies and User Policies](#) in the *Amazon S3 User Guide*.
- **Directory bucket permissions** - To grant access to this API operation, you must have the `s3express:GetInventoryConfiguration` permission in an IAM identity-based policy instead of a bucket policy. For more information about directory bucket policies and

permissions, see [AWS Identity and Access Management \(IAM\) for S3 Express One Zone](#) in the *Amazon S3 User Guide*.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `s3express-control.region-code.amazonaws.com`.

For information about the Amazon S3 inventory feature, see [Amazon S3 Inventory](#)

The following operations are related to `ListBucketInventoryConfigurations`:

- [GetBucketInventoryConfiguration](#)
- [DeleteBucketInventoryConfiguration](#)
- [PutBucketInventoryConfiguration](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?inventory&continuation-token=ContinuationToken HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

[Bucket](#)

The name of the bucket containing the inventory configurations to retrieve.

Directory buckets - When you use this operation with a directory bucket, you must use path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name`. Virtual-hosted-style requests aren't supported.

Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must also follow the format `bucket-base-name--zone-id--x-s3` (for example, `DOC-EXAMPLE-BUCKET--usw2-az1--x-s3`). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*.

Required: Yes

[continuation-token](#)

The marker used to continue an inventory configuration listing that has been truncated. Use the `NextContinuationToken` from a previously truncated list response to continue the listing. The continuation token is an opaque value that Amazon S3 understands.

[x-amz-expected-bucket-owner](#)

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Note

For directory buckets, this header is not supported in this API operation. If you specify this header, the request fails with the HTTP status code 501 Not Implemented.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<ListInventoryConfigurationsResult>
  <ContinuationToken>string</ContinuationToken>
  <InventoryConfiguration>
    <Destination>
      <S3BucketDestination>
        <AccountId>string</AccountId>
        <Bucket>string</Bucket>
        <Encryption>
          <SSE-KMS>
```

```

        <KeyId>string</KeyId>
    </SSE-KMS>
    <SSE-S3>
    </SSE-S3>
</Encryption>
    <Format>string</Format>
    <Prefix>string</Prefix>
</S3BucketDestination>
</Destination>
<Filter>
    <Prefix>string</Prefix>
</Filter>
<Id>string</Id>
<IncludedObjectVersions>string</IncludedObjectVersions>
<IsEnabled>boolean</IsEnabled>
<OptionalFields>
    <Field>string</Field>
</OptionalFields>
<Schedule>
    <Frequency>string</Frequency>
</Schedule>
</InventoryConfiguration>
...
<IsTruncated>boolean</IsTruncated>
<NextContinuationToken>string</NextContinuationToken>
</ListInventoryConfigurationsResult>

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

ListInventoryConfigurationsResult

Root level tag for the ListInventoryConfigurationsResult parameters.

Required: Yes

ContinuationToken

If sent in the request, the marker that is used as a starting point for this inventory configuration list response.

Type: String

InventoryConfiguration

The list of inventory configurations for a bucket.

Type: Array of [InventoryConfiguration](#) data types

IsTruncated

Tells whether the returned list of inventory configurations is complete. A value of true indicates that the list is not complete and the NextContinuationToken is provided for a subsequent request.

Type: Boolean

NextContinuationToken

The marker used to continue this inventory configuration listing. Use the NextContinuationToken from this response to continue the listing in a subsequent request. The continuation token is an opaque value that Amazon S3 understands.

Type: String

Examples

Sample Request

The following request returns the inventory configurations in example-bucket.

```
GET /?inventory HTTP/1.1
Host: example-bucket.s3.<Region>.amazonaws.com
x-amz-date: 20160430T233541Z
Authorization: authorization string
Content-Type: text/plain
```

Sample Response

Delete the metric configuration with a specified ID, which disables the CloudWatch metrics with the ExampleMetrics value for the FilterId dimension.

```
HTTP/1.1 200 OK
x-amz-id-2: gyB+3jRPnrkN98ZajxHXr3u7EFM67bNgSAXexeEHndCX/7GRnfTXxReKUQF28IfP
```

```
x-amz-request-id: 3B3C7C725673C630
Date: Sat, 30 Apr 2016 23:29:37 GMT
Content-Type: application/xml
Content-Length: length
Connection: close
Server: AmazonS3

<?xml version="1.0" encoding="UTF-8"?>
<ListInventoryConfigurationsResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <InventoryConfiguration>
    <Id>report1</Id>
    <IsEnabled>true</IsEnabled>
    <Destination>
      <S3BucketDestination>
        <Format>CSV</Format>
        <AccountId>123456789012</AccountId>
        <Bucket>arn:aws:s3:::destination-bucket</Bucket>
        <Prefix>prefix1</Prefix>
      </S3BucketDestination>
    </Destination>
    <Schedule>
      <Frequency>Daily</Frequency>
    </Schedule>
    <Filter>
      <Prefix>prefix/One</Prefix>
    </Filter>
    <IncludedObjectVersions>All</IncludedObjectVersions>
    <OptionalFields>
      <Field>Size</Field>
      <Field>LastModifiedDate</Field>
      <Field>ETag</Field>
      <Field>StorageClass</Field>
      <Field>IsMultipartUploaded</Field>
      <Field>ReplicationStatus</Field>
    </OptionalFields>
  </InventoryConfiguration>
  <InventoryConfiguration>
    <Id>report2</Id>
    <IsEnabled>true</IsEnabled>
    <Destination>
      <S3BucketDestination>
        <Format>CSV</Format>
        <AccountId>123456789012</AccountId>
        <Bucket>arn:aws:s3:::bucket2</Bucket>
```

```

        <Prefix>prefix2</Prefix>
    </S3BucketDestination>
</Destination>
<Schedule>
    <Frequency>Daily</Frequency>
</Schedule>
<Filter>
    <Prefix>prefix/Two</Prefix>
</Filter>
<IncludedObjectVersions>All</IncludedObjectVersions>
<OptionalFields>
    <Field>Size</Field>
    <Field>LastModifiedDate</Field>
    <Field>ETag</Field>
    <Field>StorageClass</Field>
    <Field>IsMultipartUploaded</Field>
    <Field>ReplicationStatus</Field>
    <Field>ObjectLockRetainUntilDate</Field>
    <Field>ObjectLockMode</Field>
    <Field>ObjectLockLegalHoldStatus</Field>
</OptionalFields>
</InventoryConfiguration>
<InventoryConfiguration>
    <Id>report3</Id>
    <IsEnabled>true</IsEnabled>
    <Destination>
        <S3BucketDestination>
            <Format>CSV</Format>
            <AccountId>123456789012</AccountId>
            <Bucket>arn:aws:s3:::bucket3</Bucket>
            <Prefix>prefix3</Prefix>
        </S3BucketDestination>
    </Destination>
    <Schedule>
        <Frequency>Daily</Frequency>
    </Schedule>
    <Filter>
        <Prefix>prefix/Three</Prefix>
    </Filter>
    <IncludedObjectVersions>All</IncludedObjectVersions>
    <OptionalFields>
        <Field>Size</Field>
        <Field>LastModifiedDate</Field>
        <Field>ETag</Field>

```

```
<Field>StorageClass</Field>
<Field>IsMultipartUploaded</Field>
<Field>ReplicationStatus</Field>
</OptionalFields>
</InventoryConfiguration>
...
<IsTruncated>>false</IsTruncated>
<!-- If ContinuationToken was provided in the request. -->
<ContinuationToken>...</ContinuationToken>
<!-- if IsTruncated == true -->
<IsTruncated>>true</IsTruncated>
<NextContinuationToken>...</NextContinuationToken>
</ListInventoryConfigurationsResult>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListBucketMetricsConfigurations

Service: Amazon S3

Lists the metrics configurations for the bucket. The metrics configurations are only for the request metrics of the bucket and do not provide information on daily storage metrics. You can have up to 1,000 configurations per bucket.

Note

Directory buckets - For directory buckets, you must make requests for this API operation to the Regional endpoint. These endpoints support path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name`. Virtual-hosted-style requests aren't supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

This action supports list pagination and does not return more than 100 configurations at a time. Always check the `IsTruncated` element in the response. If there are no more configurations to list, `IsTruncated` is set to false. If there are more configurations to list, `IsTruncated` is set to true, and there is a value in `NextContinuationToken`. You use the `NextContinuationToken` value to continue the pagination of the list by passing the value in `continuation-token` in the request to GET the next page.

Permissions

To use this operation, you must have permissions to perform the `s3:GetMetricsConfiguration` action. The bucket owner has this permission by default. The bucket owner can grant this permission to others. For more information about permissions, see [Permissions Related to Bucket Subresource Operations](#) and [Managing Access Permissions to Your Amazon S3 Resources](#).

- **General purpose bucket permissions** - The `s3:GetMetricsConfiguration` permission is required in a policy. For more information about general purpose buckets permissions, see [Using Bucket Policies and User Policies](#) in the *Amazon S3 User Guide*.
- **Directory bucket permissions** - To grant access to this API operation, you must have the `s3express:GetMetricsConfiguration` permission in an IAM identity-based policy

instead of a bucket policy. Cross-account access to this API operation isn't supported. This operation can only be performed by the AWS account that owns the resource. For more information about directory bucket policies and permissions, see [AWS Identity and Access Management \(IAM\) for S3 Express One Zone](#) in the *Amazon S3 User Guide*.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `s3express-control.region-code.amazonaws.com`.

For more information about metrics configurations and CloudWatch request metrics, see [Monitoring Metrics with Amazon CloudWatch](#).

The following operations are related to `ListBucketMetricsConfigurations`:

- [PutBucketMetricsConfiguration](#)
- [GetBucketMetricsConfiguration](#)
- [DeleteBucketMetricsConfiguration](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?metrics&continuation-token=ContinuationToken HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket containing the metrics configurations to retrieve.

Directory buckets - When you use this operation with a directory bucket, you must use path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name`. Virtual-hosted-style requests aren't supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must also follow the format `bucket-base-name--zone-id--x-s3` (for example, `DOC-EXAMPLE-BUCKET--usw2-az1--x-s3`). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*.

Required: Yes

continuation-token

The marker that is used to continue a metrics configuration listing that has been truncated. Use the `NextContinuationToken` from a previously truncated list response to continue the listing. The continuation token is an opaque value that Amazon S3 understands.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Note

For directory buckets, this header is not supported in this API operation. If you specify this header, the request fails with the HTTP status code 501 Not Implemented.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<ListMetricsConfigurationsResult>
  <IsTruncated>boolean</IsTruncated>
  <ContinuationToken>string</ContinuationToken>
  <NextContinuationToken>string</NextContinuationToken>
  <MetricsConfiguration>
    <Filter>
```

```

    <AccessPointArn>string</AccessPointArn>
    <And>
      <AccessPointArn>string</AccessPointArn>
      <Prefix>string</Prefix>
      <Tag>
        <Key>string</Key>
        <Value>string</Value>
      </Tag>
      ...
    </And>
    <Prefix>string</Prefix>
    <Tag>
      <Key>string</Key>
      <Value>string</Value>
    </Tag>
  </Filter>
  <Id>string</Id>
</MetricsConfiguration>
...
</ListMetricsConfigurationsResult>

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

ListMetricsConfigurationsResult

Root level tag for the ListMetricsConfigurationsResult parameters.

Required: Yes

ContinuationToken

The marker that is used as a starting point for this metrics configuration list response. This value is present if it was sent in the request.

Type: String

IsTruncated

Indicates whether the returned list of metrics configurations is complete. A value of true indicates that the list is not complete and the NextContinuationToken will be provided for a subsequent request.

Type: Boolean

MetricsConfiguration

The list of metrics configurations for a bucket.

Type: Array of [MetricsConfiguration](#) data types

NextContinuationToken

The marker used to continue a metrics configuration listing that has been truncated. Use the `NextContinuationToken` from a previously truncated list response to continue the listing. The continuation token is an opaque value that Amazon S3 understands.

Type: String

Examples

Sample Request

Delete the metric configuration with a specified ID, which disables the CloudWatch metrics with the `ExampleMetrics` value for the `FilterId` dimension.

```
GET /?metrics HTTP/1.1
Host: examplebucket.s3.<Region>.amazonaws.com
x-amz-date: Thu, 15 Nov 2016 00:17:21 GMT
Authorization: signatureValue
```

Sample Response

Delete the metric configuration with a specified ID, which disables the CloudWatch metrics with the `ExampleMetrics` value for the `FilterId` dimension.

```
HTTP/1.1 200 OK
x-amz-id-2: ITnGT1y4REXAMPLEPi4hk1TXouTf0hccUjo0iCPEXAMPLEutBj3M7fPG1W02SEWp
x-amz-request-id: 51991EXAMPLE5321
Date: Thu, 15 Nov 2016 00:17:22 GMT
Server: AmazonS3
Content-Length: 758
```

```
<?xml version="1.0" encoding="UTF-8"?>
<ListMetricsConfigurationsResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <MetricsConfiguration>
    <Id>EntireBucket</Id>
  </MetricsConfiguration>
  <MetricsConfiguration>
    <Id>Documents</Id>
    <Filter>
      <Prefix>documents/</Prefix>
    </Filter>
  </MetricsConfiguration>
  <MetricsConfiguration>
    <Id>BlueDocuments</Id>
    <Filter>
      <And>
        <Prefix>documents/</Prefix>
        <Tag>
          <Key>class</Key>
          <Value>blue</Value>
        </Tag>
      </And>
    </Filter>
  </MetricsConfiguration>
  <IsTruncated>>false</IsTruncated>
</ListMetricsConfigurationsResult>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)

- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListBuckets

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Returns a list of all buckets owned by the authenticated sender of the request. To grant IAM permission to use this operation, you must add the `s3:ListAllMyBuckets` policy action.

For information about Amazon S3 buckets, see [Creating, configuring, and working with Amazon S3 buckets](#).

Important

We strongly recommend using only paginated `ListBuckets` requests. Unpaginated `ListBuckets` requests are only supported for AWS accounts set to the default general purpose bucket quota of 10,000. If you have an approved general purpose bucket quota above 10,000, you must send paginated `ListBuckets` requests to list your account's buckets. All unpaginated `ListBuckets` requests will be rejected for AWS accounts with a general purpose bucket quota greater than 10,000.

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?bucket-region=BucketRegion&continuation-token=ContinuationToken&max-buckets=MaxBuckets&prefix=Prefix HTTP/1.1
Host: s3.amazonaws.com
```

URI Request Parameters

The request uses the following URI parameters.

bucket-region

Limits the response to buckets that are located in the specified AWS Region. The AWS Region must be expressed according to the AWS Region code, such as `us-west-2` for the US West (Oregon) Region. For a list of the valid values for all of the AWS Regions, see [Regions and Endpoints](#).

Note

Requests made to a Regional endpoint that is different from the `bucket-region` parameter are not supported. For example, if you want to limit the response to your buckets in Region `us-west-2`, the request must be made to an endpoint in Region `us-west-2`.

continuation-token

`ContinuationToken` indicates to Amazon S3 that the list is being continued on this bucket with a token. `ContinuationToken` is obfuscated and is not a real key. You can use this `ContinuationToken` for pagination of the list results.

Length Constraints: Minimum length of 0. Maximum length of 1024.

Required: No.

Note

If you specify the `bucket-region`, `prefix`, or `continuation-token` query parameters without using `max-buckets` to set the maximum number of buckets returned in the response, Amazon S3 applies a default page size of 10,000 and provides a continuation token if there are more buckets.

max-buckets

Maximum number of buckets to be returned in response. When the number is more than the count of buckets that are owned by an AWS account, return all the buckets in response.

Valid Range: Minimum value of 1. Maximum value of 10000.

[prefix](#)

Limits the response to bucket names that begin with the specified bucket name prefix.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<ListAllMyBucketsResult>
  <Buckets>
    <Bucket>
      <BucketArn>string</BucketArn>
      <BucketRegion>string</BucketRegion>
      <CreationDate>timestamp</CreationDate>
      <Name>string</Name>
    </Bucket>
  </Buckets>
  <Owner>
    <DisplayName>string</DisplayName>
    <ID>string</ID>
  </Owner>
  <ContinuationToken>string</ContinuationToken>
  <Prefix>string</Prefix>
</ListAllMyBucketsResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[ListAllMyBucketsResult](#)

Root level tag for the ListAllMyBucketsResult parameters.

Required: Yes

Buckets

The list of buckets owned by the requester.

Type: Array of [Bucket](#) data types

ContinuationToken

ContinuationToken is included in the response when there are more buckets that can be listed with pagination. The next ListBuckets request to Amazon S3 can be continued with this ContinuationToken. ContinuationToken is obfuscated and is not a real bucket.

Type: String

Owner

The owner of the buckets listed.

Type: [Owner](#) data type

Prefix

If Prefix was sent with the request, it is included in the response.

All bucket names in the response begin with the specified bucket name prefix.

Type: String

Examples

Example 1: Unpaginated ListBuckets request

This example lists all the buckets in your account in a single unpaginated response. Unpaginated requests are only supported for AWS accounts that have the default service quota of 10,000 buckets. If you have an approved general purpose bucket quota that is greater than 10,000 buckets, all unpaginated requests will be rejected for your account.

```
GET / host:s3.us-east-2.amazonaws.com HTTP/1.1
```

```
HTTP/1.1 200 OK
<ListAllMyBucketsResult>
  <Buckets>
    <Bucket>
```

```

    <CreationDate>2019-12-11T23:32:47+00:00</CreationDate>
    <Name>amzn-s3-demo-bucket</Name>
  </Bucket>
  <Bucket>
    <CreationDate>2019-11-10T23:32:13+00:00</CreationDate>
    <Name>amzn-s3-demo-bucket1</Name>
  </Bucket>
</Buckets>
<Owner>
  <ID>AIDACKCEVSQ6C2EXAMPLE</ID>
</Owner>
</ListAllMyBucketsResult>

```

Example 2: Paginated ListBuckets request

The following example request lists all buckets in your account using pagination. It gets the first page of results with the page size set to 1000 buckets. The response returns a ContinuationToken that is used in **Example 3** to list the next 1000 buckets.

```
GET /?max-buckets=1000&host:s3.us-east-2.amazonaws.com HTTP/1.1
```

```

HTTP/1.1 200 OK
<ListAllMyBucketsResult>
  <Buckets>
    <Bucket>
      <CreationDate>2024-11-14T23:32:47+00:00</CreationDate>
      <Name>amzn-s3-demo-bucket</Name>
      <BucketRegion>us-east-1</BucketRegion>
    </Bucket>
    <Bucket>
      <CreationDate>2024-11-14T23:32:13+00:00</CreationDate>
      <Name>amzn-s3-demo-bucket1</Name>
      <BucketRegion>us-east-2</BucketRegion>
    </Bucket>
  </Buckets>
  <Owner>
    <ID>AIDACKCEVSQ6C2EXAMPLE</ID>
  </Owner>
  <ContinuationToken>eyJNYXJrZXIiOiBudWxsLCAiYm90b190cnVuY2F0ZV9hbW91bnQiOiAxZQ==</
ContinuationToken>

```



```
</ListAllMyBucketsResult>
```

Example 3: Paginated ListBuckets request with continuation token

This example request uses the token returned in **Example 2** to return the next 1000 buckets. Continue until there are no more results. If you do not receive a continuation token with your initial paginated ListBuckets request, then your single paginated request returned all of the buckets in your account.

```
GET /?max-buckets=1000&continuation-  
token=eyJNYXJrZXIiOiBudWxsLCAiYm90b190cnVuY2F0ZV9hbW91bnQiOiAxZjQ== host:s3.us-  
east-2.amazonaws.com HTTP/1.1
```

```
HTTP/1.1 200 OK  
<ListAllMyBucketsResult>  
  <Buckets>  
    <Bucket>  
      <CreationDate>2024-11-14T23:32:47+00:00</CreationDate>  
      <Name>amzn-s3-demo-bucket</Name>  
      <BucketRegion>us-east-1</BucketRegion>  
    </Bucket>  
    <Bucket>  
      <CreationDate>2024-11-14T23:32:13+00:00</CreationDate>  
      <Name>amzn-s3-demo-bucket1</Name>  
      <BucketRegion>us-east-2</BucketRegion>  
    </Bucket>  
  </Buckets>  
  <Owner>  
    <ID>AIDACKCEVSQ6C2EXAMPLE</ID>  
  </Owner>  
  
  <ContinuationToken>eyJ0ZXh0VG9rZW4iOiBudWxsLCAiYm90b190cnVuY2F0ZV9hbW91bnQiEXAMPLE=</  
ContinuationToken>  
</ListAllMyBucketsResult>
```

Example 4: Paginated ListBuckets request for buckets in US East (Ohio) (us-east-2)

The following example lists all the buckets in your account in the us-east-2 Region. The first paginated response will return up to 1000 buckets. Requests made to a Regional endpoint that is different from the bucket-region parameter are not supported.

```
GET /?bucket-region=us-east-2&max-buckets=1000 host:s3.us-east-2.amazonaws.com HTTP/1.1
```

```
HTTP/1.1 200 OK
<ListAllMyBucketsResult>
  <Buckets>
    <Bucket>
      <CreationDate>2024-11-14T23:32:47+00:00</CreationDate>
      <Name>DOC-EXAMPLE-BUCKET</Name>
      <BucketRegion>us-east-2</BucketRegion>
    </Bucket>
    <Bucket>
      <CreationDate>2024-11-14T23:32:13+00:00</CreationDate>
      <Name>DOC-EXAMPLE-BUCKET1002</Name>
      <BucketRegion>us-east-2</BucketRegion>
    </Bucket>
  </Buckets>
  <Owner>
    <ID>AIDACKCEVSQ6C2EXAMPLE</ID>
  </Owner>
  <ContinuationToken>eyJ0ZXh0VG9rZW4iOiBudWxsLCAiYm90b190cnVuY2F0ZV9hbW91bnQiEXAMPLEcd
=</ContinuationToken>
</ListAllMyBucketsResult>
```

Example 5: Paginated ListBuckets request for buckets in your account that begin with amzn-s3-demo-bucket in US East (Ohio) (us-east-2)

The following example lists all the buckets in your account located in the us-east-2 Region that begin with the amzn-s3-demo-bucket bucket prefix. This request uses pagination.

```
GET /?bucket-region=us-east-2&max-buckets=1000&prefix=amzn-s3-demo-bucket host:s3.us-east-2.amazonaws.com HTTP/1.1
```

```
HTTP/1.1 200 OK
<ListAllMyBucketsResult>
  <Buckets>
    <Bucket>
      <CreationDate>2024-11-14T23:32:47+00:00</CreationDate>
      <Name>amzn-s3-demo-bucket</Name>
      <BucketRegion>us-east-2</BucketRegion>
    </Bucket>
  </Buckets>
  <Owner>
    <ID>AIDACKCEVSQ6C2EXAMPLE</ID>
  </Owner>
  <Prefix>
    amzn-s3-demo-bucket
  </Prefix>

  <ContinuationToken>eyJ0ZXh0VG9rZW4iOiBudWxsLCAiYm90b190cnVuY2F0ZV9hbW91bnQiEXAMPLE=</ContinuationToken>
</ListAllMyBucketsResult>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListDirectoryBuckets

Service: Amazon S3

Returns a list of all Amazon S3 directory buckets owned by the authenticated sender of the request. For more information about directory buckets, see [Directory buckets](#) in the *Amazon S3 User Guide*.

Note

Directory buckets - For directory buckets, you must make requests for this API operation to the Regional endpoint. These endpoints support path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name` . Virtual-hosted-style requests aren't supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

Permissions

You must have the `s3express:ListAllMyDirectoryBuckets` permission in an IAM identity-based policy instead of a bucket policy. Cross-account access to this API operation isn't supported. This operation can only be performed by the AWS account that owns the resource. For more information about directory bucket policies and permissions, see [AWS Identity and Access Management \(IAM\) for S3 Express One Zone](#) in the *Amazon S3 User Guide*.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `s3express-control.region.amazonaws.com`.

Note

The `BucketRegion` response element is not part of the `ListDirectoryBuckets` Response Syntax.

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?continuation-token=ContinuationToken&max-directory-buckets=MaxDirectoryBuckets
HTTP/1.1
Host: s3.amazonaws.com
```

URI Request Parameters

The request uses the following URI parameters.

continuation-token

`ContinuationToken` indicates to Amazon S3 that the list is being continued on buckets in this account with a token. `ContinuationToken` is obfuscated and is not a real bucket name. You can use this `ContinuationToken` for the pagination of the list results.

Length Constraints: Minimum length of 0. Maximum length of 1024.

max-directory-buckets

Maximum number of buckets to be returned in response. When the number is more than the count of buckets that are owned by an AWS account, return all the buckets in response.

Valid Range: Minimum value of 0. Maximum value of 1000.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
```

```
<ListAllMyDirectoryBucketsResult>
  <Buckets>
    <Bucket>
      <BucketArn>string</BucketArn>
      <BucketRegion>string</BucketRegion>
      <CreationDate>timestamp</CreationDate>
      <Name>string</Name>
    </Bucket>
  </Buckets>
  <ContinuationToken>string</ContinuationToken>
</ListAllMyDirectoryBucketsResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[ListAllMyDirectoryBucketsResult](#)

Root level tag for the ListAllMyDirectoryBucketsResult parameters.

Required: Yes

[Buckets](#)

The list of buckets owned by the requester.

Type: Array of [Bucket](#) data types

[ContinuationToken](#)

If ContinuationToken was sent with the request, it is included in the response. You can use the returned ContinuationToken for pagination of the list response.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 1024.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListMultipartUploads

Service: Amazon S3

This operation lists in-progress multipart uploads in a bucket. An in-progress multipart upload is a multipart upload that has been initiated by the `CreateMultipartUpload` request, but has not yet been completed or aborted.

Note

Directory buckets - If multipart uploads in a directory bucket are in progress, you can't delete the bucket until all the in-progress multipart uploads are aborted or completed. To delete these in-progress multipart uploads, use the `ListMultipartUploads` operation to list the in-progress multipart uploads in the bucket and use the `AbortMultipartUpload` operation to abort all the in-progress multipart uploads.

The `ListMultipartUploads` operation returns a maximum of 1,000 multipart uploads in the response. The limit of 1,000 multipart uploads is also the default value. You can further limit the number of uploads in a response by specifying the `max-uploads` request parameter. If there are more than 1,000 multipart uploads that satisfy your `ListMultipartUploads` request, the response returns an `IsTruncated` element with the value of `true`, a `NextKeyMarker` element, and a `NextUploadIdMarker` element. To list the remaining multipart uploads, you need to make subsequent `ListMultipartUploads` requests. In these requests, include two query parameters: `key-marker` and `upload-id-marker`. Set the value of `key-marker` to the `NextKeyMarker` value from the previous response. Similarly, set the value of `upload-id-marker` to the `NextUploadIdMarker` value from the previous response.

Note

Directory buckets - The `upload-id-marker` element and the `NextUploadIdMarker` element aren't supported by directory buckets. To list the additional multipart uploads, you only need to set the value of `key-marker` to the `NextKeyMarker` value from the previous response.

For more information about multipart uploads, see [Uploading Objects Using Multipart Upload](#) in the *Amazon S3 User Guide*.

Note

Directory buckets - For directory buckets, you must make requests for this API operation to the Zonal endpoint. These endpoints support virtual-hosted-style requests in the format `https://amzn-s3-demo-bucket.s3express-zone-id.region-code.amazonaws.com/key-name` . Path-style requests are not supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

Permissions

- **General purpose bucket permissions** - For information about permissions required to use the multipart upload API, see [Multipart Upload and Permissions](#) in the *Amazon S3 User Guide*.
- **Directory bucket permissions** - To grant access to this API operation on a directory bucket, we recommend that you use the [CreateSession](#) API operation for session-based authorization. Specifically, you grant the `s3express:CreateSession` permission to the directory bucket in a bucket policy or an IAM identity-based policy. Then, you make the `CreateSession` API call on the bucket to obtain a session token. With the session token in your request header, you can make API requests to this operation. After the session token expires, you make another `CreateSession` API call to generate a new session token for use. AWS CLI or SDKs create session and refresh the session token automatically to avoid service interruptions when a session expires. For more information about authorization, see [CreateSession](#).

Sorting of multipart uploads in response

- **General purpose bucket** - In the `ListMultipartUploads` response, the multipart uploads are sorted based on two criteria:
 - **Key-based sorting** - Multipart uploads are initially sorted in ascending order based on their object keys.
 - **Time-based sorting** - For uploads that share the same object key, they are further sorted in ascending order based on the upload initiation time. Among uploads with the same key, the one that was initiated first will appear before the ones that were initiated later.
- **Directory bucket** - In the `ListMultipartUploads` response, the multipart uploads aren't sorted lexicographically based on the object keys.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `Bucket-name.s3express-zone-id.region-code.amazonaws.com`.

The following operations are related to `ListMultipartUploads`:

- [CreateMultipartUpload](#)
- [UploadPart](#)
- [CompleteMultipartUpload](#)
- [ListParts](#)
- [AbortMultipartUpload](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?uploads&delimiter=Delimiter&encoding-type=EncodingType&key-marker=KeyMarker&max-uploads=MaxUploads&prefix=Prefix&upload-id-marker=UploadIdMarker HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
x-amz-request-payer: RequestPayer
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket to which the multipart upload was initiated.

Directory buckets - When you use this operation with a directory bucket, you must use virtual-hosted-style requests in the format `Bucket-name.s3express-zone-id.region-`

`code.amazonaws.com`. Path-style requests are not supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must follow the format `bucket-base-name--zone-id--x-s3` (for example, `amzn-s3-demo-bucket--usw2-az1--x-s3`). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form `AccessPointName-AccountId.s3-accesspoint.Region.amazonaws.com`. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

 Note

Object Lambda access points are not supported by directory buckets.

S3 on Outposts - When you use this action with S3 on Outposts, you must direct requests to the S3 on Outposts hostname. The S3 on Outposts hostname takes the form `AccessPointName-AccountId.outpostID.s3-outposts.Region.amazonaws.com`. When you use this action with S3 on Outposts, the destination bucket must be the Outposts access point ARN or the access point alias. For more information about S3 on Outposts, see [What is S3 on Outposts?](#) in the *Amazon S3 User Guide*.

Required: Yes

delimiter

Character you use to group keys.

All keys that contain the same string between the prefix, if specified, and the first occurrence of the delimiter after the prefix are grouped under a single result element, `CommonPrefixes`. If you don't specify the prefix parameter, then the substring starts at the beginning of the key. The keys that are grouped under `CommonPrefixes` result element are not returned elsewhere in the response.

`CommonPrefixes` is filtered out from results if it is not lexicographically greater than the `key-marker`.

 Note

Directory buckets - For directory buckets, `/` is the only supported delimiter.

encoding-type

Encoding type used by Amazon S3 to encode the [object keys](#) in the response. Responses are encoded only in UTF-8. An object key can contain any Unicode character. However, the XML 1.0 parser can't parse certain characters, such as characters with an ASCII value from 0 to 10. For characters that aren't supported in XML 1.0, you can add this parameter to request that Amazon S3 encode the keys in the response. For more information about characters to avoid in object key names, see [Object key naming guidelines](#).

 Note

When using the URL encoding type, non-ASCII characters that are used in an object's key name will be percent-encoded according to UTF-8 code values. For example, the object `test_file(3).png` will appear as `test_file%283%29.png`.

Valid Values: `url`

key-marker

Specifies the multipart upload after which listing should begin.

 Note

- **General purpose buckets** - For general purpose buckets, `key-marker` is an object key. Together with `upload-id-marker`, this parameter specifies the multipart upload after which listing should begin.

If `upload-id-marker` is not specified, only the keys lexicographically greater than the specified `key-marker` will be included in the list.

If `upload-id-marker` is specified, any multipart uploads for a key equal to the `key-marker` might also be included, provided those multipart uploads have upload IDs lexicographically greater than the specified `upload-id-marker`.

- **Directory buckets** - For directory buckets, `key-marker` is obfuscated and isn't a real object key. The `upload-id-marker` parameter isn't supported by directory buckets. To list the additional multipart uploads, you only need to set the value of `key-marker` to the `NextKeyMarker` value from the previous response.

In the `ListMultipartUploads` response, the multipart uploads aren't sorted lexicographically based on the object keys.

max-uploads

Sets the maximum number of multipart uploads, from 1 to 1,000, to return in the response body. 1,000 is the maximum number of uploads that can be returned in a response.

prefix

Lists in-progress uploads only for those keys that begin with the specified prefix. You can use prefixes to separate a bucket into different grouping of keys. (You can think of using `prefix` to make groups in the same way that you'd use a folder in a file system.)

Note

Directory buckets - For directory buckets, only prefixes that end in a delimiter (/) are supported.

upload-id-marker

Together with `key-marker`, specifies the multipart upload after which listing should begin. If `key-marker` is not specified, the `upload-id-marker` parameter is ignored. Otherwise, any multipart uploads for a key equal to the `key-marker` might be included in the list only if they have an upload ID lexicographically greater than the specified `upload-id-marker`.

Note

This functionality is not supported for directory buckets.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
x-amz-request-charged: RequestCharged
<?xml version="1.0" encoding="UTF-8"?>
<ListMultipartUploadsResult>
  <Bucket>string</Bucket>
  <KeyMarker>string</KeyMarker>
  <UploadIdMarker>string</UploadIdMarker>
```

```

<NextKeyMarker>string</NextKeyMarker>
<Prefix>string</Prefix>
<Delimiter>string</Delimiter>
<NextUploadIdMarker>string</NextUploadIdMarker>
<MaxUploads>integer</MaxUploads>
<IsTruncated>boolean</IsTruncated>
<Upload>
  <ChecksumAlgorithm>string</ChecksumAlgorithm>
  <ChecksumType>string</ChecksumType>
  <Initiated>timestamp</Initiated>
  <Initiator>
    <DisplayName>string</DisplayName>
    <ID>string</ID>
  </Initiator>
  <Key>string</Key>
  <Owner>
    <DisplayName>string</DisplayName>
    <ID>string</ID>
  </Owner>
  <StorageClass>string</StorageClass>
  <UploadId>string</UploadId>
</Upload>
...
<CommonPrefixes>
  <Prefix>string</Prefix>
</CommonPrefixes>
...
<EncodingType>string</EncodingType>
</ListMultipartUploadsResult>

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: `requester`

The following data is returned in XML format by the service.

ListMultipartUploadsResult

Root level tag for the ListMultipartUploadsResult parameters.

Required: Yes

Bucket

The name of the bucket to which the multipart upload was initiated. Does not return the access point ARN or access point alias if used.

Type: String

CommonPrefixes

If you specify a delimiter in the request, then the result returns each distinct key prefix containing the delimiter in a `CommonPrefixes` element. The distinct key prefixes are returned in the `Prefix` child element.

Note

Directory buckets - For directory buckets, only prefixes that end in a delimiter (/) are supported.

Type: Array of [CommonPrefix](#) data types

Delimiter

Contains the delimiter you specified in the request. If you don't specify a delimiter in your request, this element is absent from the response.

 Note

Directory buckets - For directory buckets, / is the only supported delimiter.

Type: String

EncodingType

Encoding type used by Amazon S3 to encode object keys in the response.

If you specify the `encoding-type` request parameter, Amazon S3 includes this element in the response, and returns encoded key name values in the following response elements:

Delimiter, KeyMarker, Prefix, NextKeyMarker, Key.

Type: String

Valid Values: url

IsTruncated

Indicates whether the returned list of multipart uploads is truncated. A value of true indicates that the list was truncated. The list can be truncated if the number of multipart uploads exceeds the limit allowed or specified by max uploads.

Type: Boolean

KeyMarker

The key at or after which the listing began.

Type: String

MaxUploads

Maximum number of multipart uploads that could have been included in the response.

Type: Integer

NextKeyMarker

When a list is truncated, this element specifies the value that should be used for the key-marker request parameter in a subsequent request.

Type: String

NextUploadIdMarker

When a list is truncated, this element specifies the value that should be used for the `upload-id-marker` request parameter in a subsequent request.

Note

This functionality is not supported for directory buckets.

Type: String

Prefix

When a prefix is provided in the request, this field contains the specified prefix. The result contains only keys starting with the specified prefix.

Note

Directory buckets - For directory buckets, only prefixes that end in a delimiter (/) are supported.

Type: String

Upload

Container for elements related to a particular multipart upload. A response can contain zero or more `Upload` elements.

Type: Array of [MultipartUpload](#) data types

UploadIdMarker

Together with `key-marker`, specifies the multipart upload after which listing should begin. If `key-marker` is not specified, the `upload-id-marker` parameter is ignored. Otherwise, any multipart uploads for a key equal to the `key-marker` might be included in the list only if they have an upload ID lexicographically greater than the specified `upload-id-marker`.

Note

This functionality is not supported for directory buckets.

Type: String

Examples

Sample Request for general purpose buckets

The following request lists three multipart uploads. The request specifies the `max-uploads` request parameter to set the maximum number of multipart uploads to return in the response body.

```
GET /?uploads&max-uploads=3 HTTP/1.1
Host: example-bucket.s3.<Region>.amazonaws.com
Date: Mon, 1 Nov 2010 20:34:56 GMT
Authorization: authorization string
```

Sample Response for general purpose buckets

The following sample response indicates that the multipart upload list was truncated and provides the `NextKeyMarker` and the `NextUploadIdMarker` elements. You specify these values in your subsequent requests to read the next set of multipart uploads. That is, send a subsequent request specifying `key-marker=my-movie2.m2ts` (value of the `NextKeyMarker` element) and `upload-id-marker=YW55IGlkZWVsdmluZydzIHVwbG9hZCBmYW1sZWQ` (value of the `NextUploadIdMarker`).

The sample response also shows a case of two multipart uploads in progress with the same key (`my-movie.m2ts`). That is, the response shows two uploads with the same key. This response shows the uploads sorted by key, and within each key the uploads are sorted in ascending order by the time the multipart upload was initiated.

```
HTTP/1.1 200 OK
x-amz-id-2: Uuag1LuByRx9e6j50nimru9p04ZVKnJ2Qz7/C1NPcfTWAtRPfTa0Fg==
x-amz-request-id: 656c76696e6727732072657175657374
Date: Mon, 1 Nov 2010 20:34:56 GMT
Content-Length: 1330
Connection: keep-alive
Server: AmazonS3

<?xml version="1.0" encoding="UTF-8"?>
```

```

<ListMultipartUploadsResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Bucket>bucket</Bucket>
  <KeyMarker></KeyMarker>
  <UploadIdMarker></UploadIdMarker>
  <NextKeyMarker>my-movie.m2ts</NextKeyMarker>
  <NextUploadIdMarker>YW55IGlkZWEd2h5IGVsdmluZydzIHVwbG9hZCBmYWlsZWQ</
NextUploadIdMarker>
  <MaxUploads>3</MaxUploads>
  <IsTruncated>true</IsTruncated>
  <Upload>
    <Key>my-divisor</Key>
    <UploadId>XMgbGlrZSB1bHZpbmcncyBub3QgaGF2aW5nIG11Y2ggbHVjaw</UploadId>
    <Initiator>
      <ID>arn:aws:iam::111122223333:user/user1-11111a31-17b5-4fb7-9df5-b11111f13de</
ID>

    </Initiator>
    <Owner>
      <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeebf76c078efc7c6caea54ba06a</ID>

    </Owner>
    <StorageClass>STANDARD</StorageClass>
    <Initiated>2010-11-10T20:48:33.000Z</Initiated>
  </Upload>
  <Upload>
    <Key>my-movie.m2ts</Key>
    <UploadId>VXBsb2FkIElEIGZvciB1bHZpbmcncyBteS1tb3ZpZS5tMnRzIHVwbG9hZA</UploadId>
    <Initiator>
      <ID>b1d16700c70b0b05597d7acd6a3f92be</ID>

    </Initiator>
    <Owner>
      <ID>b1d16700c70b0b05597d7acd6a3f92be</ID>

    </Owner>
    <StorageClass>STANDARD</StorageClass>
    <Initiated>2010-11-10T20:48:33.000Z</Initiated>
  </Upload>
  <Upload>
    <Key>my-movie.m2ts</Key>
    <UploadId>YW55IGlkZWEd2h5IGVsdmluZydzIHVwbG9hZCBmYWlsZWQ</UploadId>
    <Initiator>
      <ID>arn:aws:iam::444455556666:user/user1-22222a31-17b5-4fb7-9df5-b22222f13de</
ID>

```

```
</Initiator>
<Owner>
  <ID>b1d16700c70b0b05597d7acd6a3f92be</ID>

</Owner>
<StorageClass>STANDARD</StorageClass>
<Initiated>2010-11-10T20:49:33.000Z</Initiated>
</Upload>
</ListMultipartUploadsResult>
```

Sample Request for general purpose buckets: Using the delimiter and the prefix parameters

Assume you have a multipart upload in progress for the following keys in your bucket, example-bucket.

- photos/2006/January/sample.jpg
- photos/2006/February/sample.jpg
- photos/2006/March/sample.jpg
- videos/2006/March/sample.wmv
- sample.jpg

The following list multipart upload request specifies the delimiter parameter with value "/".

```
GET /?uploads&delimiter=/ HTTP/1.1
Host: example-bucket.s3.<Region>.amazonaws.com
Date: Mon, 1 Nov 2010 20:34:56 GMT
Authorization: authorization string
```

Sample Response for general purpose buckets

The following sample response lists multipart uploads on the specified bucket, example-bucket.

The response returns multipart upload for the sample.jpg key in an <Upload> element.

However, because all the other keys contain the specified delimiter, a distinct substring, from the beginning of the key to the first occurrence of the delimiter, from each of these keys is

returned in a `<CommonPrefixes>` element. The key substrings, `photos/` and `videos/` in the `<CommonPrefixes>` element, indicate that there are one or more in-progress multipart uploads with these key prefixes.

This is a useful scenario if you use key prefixes for your objects to create a logical folder like structure. In this case, you can interpret the result as the folders `photos/` and `videos/` have one or more multipart uploads in progress.

```
<ListMultipartUploadsResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Bucket>example-bucket</Bucket>
  <KeyMarker/>
  <UploadIdMarker/>
  <NextKeyMarker>sample.jpg</NextKeyMarker>

  <NextUploadIdMarker>Agw4MJT6ZPAVxpY0SAuGN7q4uWJJM22ZYg1N99trdp4tp088.PT6.Mh00w2E17eutfAvQfQWoa
</NextUploadIdMarker>
  <Delimiter>/</Delimiter>
  <Prefix/>
  <MaxUploads>1000</MaxUploads>
  <IsTruncated>false</IsTruncated>
  <Upload>
    <Key>sample.jpg</Key>

    <UploadId>Agw4MJT6ZPAVxpY0SAuGN7q4uWJJM22ZYg1N99trdp4tp088.PT6.Mh00w2E17eutfAvQfQWoaJgE_W2gpcx
  </UploadId>
    <Initiator>
      <ID>314133b66967d86f031c7249d1d9a80249109428335cd0ef1cdc487b4566cb1b</ID>

    </Initiator>
    <Owner>
      <ID>314133b66967d86f031c7249d1d9a80249109428335cd0ef1cdc487b4566cb1b</ID>
    </Owner>
    <StorageClass>STANDARD</StorageClass>
    <Initiated>2010-11-26T19:24:17.000Z</Initiated>
  </Upload>
  <CommonPrefixes>
    <Prefix>photos/</Prefix>
  </CommonPrefixes>
  <CommonPrefixes>
    <Prefix>videos/</Prefix>
  </CommonPrefixes>
</ListMultipartUploadsResult>
```

Sample Request for general purpose buckets

In addition to the delimiter parameter, you can filter results by adding a prefix parameter as shown in the following request.

```
GET /?uploads&delimiter=/&prefix=photos/2006/ HTTP/1.1
Host: example-bucket.s3.<Region>.amazonaws.com
Date: Mon, 1 Nov 2010 20:34:56 GMT
Authorization: authorization string
```

Sample Response for general purpose buckets

In this case, the response will include only multipart uploads for keys that start with the specified prefix. The value returned in the <CommonPrefixes> element is a substring from the beginning of the key to the first occurrence of the specified delimiter after the prefix.

```
<?xml version="1.0" encoding="UTF-8"?>
<ListMultipartUploadsResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Bucket>example-bucket</Bucket>
  <KeyMarker/>
  <UploadIdMarker/>
  <NextKeyMarker/>
  <NextUploadIdMarker/>
  <Delimiter>/</Delimiter>
  <Prefix>photos/2006/</Prefix>
  <MaxUploads>1000</MaxUploads>
  <IsTruncated>false</IsTruncated>
  <CommonPrefixes>
    <Prefix>photos/2006/February/</Prefix>
  </CommonPrefixes>
  <CommonPrefixes>
    <Prefix>photos/2006/January/</Prefix>
  </CommonPrefixes>
  <CommonPrefixes>
    <Prefix>photos/2006/March/</Prefix>
  </CommonPrefixes>
</ListMultipartUploadsResult>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListObjectAnnotations

Service: Amazon S3

Lists the annotations attached to an Amazon S3 object. Results are paginated, with a maximum of 1,000 annotations per object. Use the `AnnotationPrefix` parameter to filter the results by name prefix.

To use this operation, you must have the `s3:ListObjectAnnotations` permission.

Note

Annotations are not supported by the following features: S3 Inventory Reports, API Gateway, S3 Storage Lens, Amazon S3 File Gateway, Amazon FSx, S3 on Outposts, and S3 Express One Zone (directory buckets).

The following operations are related to `ListObjectAnnotations`:

- [PutObjectAnnotation](#)
- [GetObjectAnnotation](#)
- [DeleteObjectAnnotation](#)

Request Syntax

```
GET /{Key+}?annotation&annotation-prefix=AnnotationPrefix&continuation-  
token=ContinuationToken&max-annotation-results=MaxAnnotationResults&versionId=VersionId  
HTTP/1.1  
Host: Bucket.s3.amazonaws.com  
x-amz-request-payer: RequestPayer  
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

[annotation-prefix](#)

Filter results to annotations whose name begins with the specified prefix.

Bucket

The name of the bucket that contains the object.

Required: Yes

continuation-token

Continuation token returned by a previous request to retrieve the next page.

Key

The object key.

Length Constraints: Minimum length of 1.

Required: Yes

max-annotation-results

The maximum number of annotations to return in the response. Maximum is 1,000.

Valid Range: Minimum value of 1. Maximum value of 1000.

versionId

The version ID of the object.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner.

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
x-amz-object-version-id: ObjectVersionId
x-amz-request-charged: RequestCharged
<?xml version="1.0" encoding="UTF-8"?>
<ListObjectAnnotationsOutput>
  <Annotations>
    <AnnotationEntry>
      <AnnotationName>string</AnnotationName>
      <ChecksumAlgorithm>string</ChecksumAlgorithm>
      ...
      <ETag>string</ETag>
      <LastModified>timestamp</LastModified>
      <ReplicationStatus>string</ReplicationStatus>
      <Size>long</Size>
    </AnnotationEntry>
  </Annotations>
  <Bucket>string</Bucket>
  <Key>string</Key>
  <AnnotationPrefix>string</AnnotationPrefix>
  <MaxAnnotationResults>integer</MaxAnnotationResults>
  <AnnotationCount>integer</AnnotationCount>
  <ContinuationToken>string</ContinuationToken>
  <NextContinuationToken>string</NextContinuationToken>
</ListObjectAnnotationsOutput>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

[x-amz-object-version-id](#)

The version ID of the object.

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

The following data is returned in XML format by the service.

ListObjectAnnotationsOutput

Root level tag for the ListObjectAnnotationsOutput parameters.

Required: Yes

AnnotationCount

The number of annotations returned.

Type: Integer

AnnotationPrefix

The prefix used to filter the response.

Type: String

Annotations

The list of annotations attached to the object.

Type: Array of [AnnotationEntry](#) data types

Bucket

The bucket name.

Type: String

ContinuationToken

The continuation token used in this request.

Type: String

Key

The object key.

Type: String

Length Constraints: Minimum length of 1.

MaxAnnotationResults

The maximum number of annotations returned in the response.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 1000.

NextContinuationToken

The continuation token to use to retrieve the next page of results.

Type: String

Errors

InvalidPrefix

The annotation prefix you provided is invalid.

HTTP Status Code: 400

NoSuchBucket

The specified bucket does not exist.

HTTP Status Code: 404

NoSuchKey

The specified key does not exist.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListObjects

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Returns some or all (up to 1,000) of the objects in a bucket. You can use the request parameters as selection criteria to return a subset of the objects in a bucket. A 200 OK response can contain valid or invalid XML. Be sure to design your application to parse the contents of the response and handle it appropriately.

Important

This action has been revised. We recommend that you use the newer version, [ListObjectsV2](#), when developing applications. For backward compatibility, Amazon S3 continues to support ListObjects.

The following operations are related to ListObjects:

- [ListObjectsV2](#)
- [GetObject](#)
- [PutObject](#)
- [CreateBucket](#)
- [ListBuckets](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?delimiter=Delimiter&encoding-type=EncodingType&marker=Marker&max-  
keys=MaxKeys&prefix=Prefix HTTP/1.1  
Host: Bucket.s3.amazonaws.com  
x-amz-request-payer: RequestPayer  
x-amz-expected-bucket-owner: ExpectedBucketOwner  
x-amz-optional-object-attributes: OptionalObjectAttributes
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket containing the objects.

Directory buckets - When you use this operation with a directory bucket, you must use virtual-hosted-style requests in the format *Bucket-name.s3express-zone-id.region-code.amazonaws.com*. Path-style requests are not supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must follow the format *bucket-base-name--zone-id--x-s3* (for example, *amzn-s3-demo-bucket--usw2-az1--x-s3*). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form *AccessPointName-AccountId.s3-accesspoint.Region.amazonaws.com*. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

Note

Object Lambda access points are not supported by directory buckets.

S3 on Outposts - When you use this action with S3 on Outposts, you must direct requests to the S3 on Outposts hostname. The S3 on Outposts hostname takes the form *AccessPointName-AccountId.outpostID.s3-outposts.Region.amazonaws.com*. When you use this action with S3 on Outposts, the destination bucket must be the Outposts access point ARN or the access point alias. For more information about S3 on Outposts, see [What is S3 on Outposts?](#) in the *Amazon S3 User Guide*.

Required: Yes

delimiter

A delimiter is a character that you use to group keys.

CommonPrefixes is filtered out from results if it is not lexicographically greater than the key-marker.

encoding-type

Encoding type used by Amazon S3 to encode the [object keys](#) in the response. Responses are encoded only in UTF-8. An object key can contain any Unicode character. However, the XML 1.0 parser can't parse certain characters, such as characters with an ASCII value from 0 to 10. For characters that aren't supported in XML 1.0, you can add this parameter to request that Amazon S3 encode the keys in the response. For more information about characters to avoid in object key names, see [Object key naming guidelines](#).

Note

When using the URL encoding type, non-ASCII characters that are used in an object's key name will be percent-encoded according to UTF-8 code values. For example, the object `test_file(3).png` will appear as `test_file%283%29.png`.

Valid Values: url

marker

Marker is where you want Amazon S3 to start listing from. Amazon S3 starts listing after this specified key. Marker can be any key in the bucket.

max-keys

Sets the maximum number of keys returned in the response. By default, the action returns up to 1,000 key names. The response might contain fewer keys but will never contain more.

prefix

Limits the response to keys that begin with the specified prefix.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-optional-object-attributes

Specifies the optional fields that you want returned in the response. Fields that you do not specify are not returned.

Valid Values: RestoreStatus

x-amz-request-payer

Confirms that the requester knows that she or he will be charged for the list objects request. Bucket owners need not specify this parameter in their requests.

Valid Values: requester

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
x-amz-request-charged: RequestCharged
<?xml version="1.0" encoding="UTF-8"?>
<ListBucketResult>
  <IsTruncated>boolean</IsTruncated>
  <Marker>string</Marker>
  <NextMarker>string</NextMarker>
  <Contents>
    <ChecksumAlgorithm>string</ChecksumAlgorithm>
    ...
    <ChecksumType>string</ChecksumType>
    <ETag>string</ETag>
    <Key>string</Key>
```

```

<LastModified>timestamp</LastModified>
<Owner>
  <DisplayName>string</DisplayName>
  <ID>string</ID>
</Owner>
<RestoreStatus>
  <IsRestoreInProgress>boolean</IsRestoreInProgress>
  <RestoreExpiryDate>timestamp</RestoreExpiryDate>
</RestoreStatus>
<Size>long</Size>
<StorageClass>string</StorageClass>
</Contents>
...
<Name>string</Name>
<Prefix>string</Prefix>
<Delimiter>string</Delimiter>
<MaxKeys>integer</MaxKeys>
<CommonPrefixes>
  <Prefix>string</Prefix>
</CommonPrefixes>
...
<EncodingType>string</EncodingType>
</ListBucketResult>

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

The following data is returned in XML format by the service.

ListBucketResult

Root level tag for the ListBucketResult parameters.

Required: Yes

CommonPrefixes

All of the keys (up to 1,000) rolled up in a common prefix count as a single return when calculating the number of returns.

A response can contain CommonPrefixes only if you specify a delimiter.

CommonPrefixes contains all (if there are any) keys between Prefix and the next occurrence of the string specified by the delimiter.

CommonPrefixes lists keys that act like subdirectories in the directory specified by Prefix.

For example, if the prefix is notes/ and the delimiter is a slash (/), as in notes/summer/july, the common prefix is notes/summer/. All of the keys that roll up into a common prefix count as a single return when calculating the number of returns.

Type: Array of [CommonPrefix](#) data types

Contents

Metadata about each object returned.

Type: Array of [Object](#) data types

Delimiter

Causes keys that contain the same string between the prefix and the first occurrence of the delimiter to be rolled up into a single result element in the CommonPrefixes collection. These rolled-up keys are not returned elsewhere in the response. Each rolled-up result counts as only one return against the MaxKeys value.

Type: String

EncodingType

Encoding type used by Amazon S3 to encode the [object keys](#) in the response. Responses are encoded only in UTF-8. An object key can contain any Unicode character. However, the XML 1.0 parser can't parse certain characters, such as characters with an ASCII value from 0 to 10. For

characters that aren't supported in XML 1.0, you can add this parameter to request that Amazon S3 encode the keys in the response. For more information about characters to avoid in object key names, see [Object key naming guidelines](#).

Note

When using the URL encoding type, non-ASCII characters that are used in an object's key name will be percent-encoded according to UTF-8 code values. For example, the object `test_file(3).png` will appear as `test_file%283%29.png`.

Type: String

Valid Values: `url`

IsTruncated

A flag that indicates whether Amazon S3 returned all of the results that satisfied the search criteria.

Type: Boolean

Marker

Indicates where in the bucket listing begins. Marker is included in the response if it was sent with the request.

Type: String

MaxKeys

The maximum number of keys returned in the response body.

Type: Integer

Name

The bucket name.

Type: String

NextMarker

When the response is truncated (the `IsTruncated` element value in the response is `true`), you can use the key name in this field as the `marker` parameter in the subsequent request to get the next set of objects. Amazon S3 lists objects in alphabetical order.

Note

This element is returned only if you have the `delimiter` request parameter specified. If the response does not include the `NextMarker` element and it is truncated, you can use the value of the last `Key` element in the response as the `marker` parameter in the subsequent request to get the next set of object keys.

Type: String

Prefix

Keys that begin with the indicated prefix.

Type: String

Errors**NoSuchBucket**

The specified bucket does not exist.

HTTP Status Code: 404

Examples**Sample Request**

This request returns the objects in `BucketName`.

```
GET / HTTP/1.1
Host: BucketName.s3.<Region>.amazonaws.com
Date: Wed, 12 Oct 2009 17:50:00 GMT
Authorization: authorization string
Content-Type: text/plain
```

Sample Response

This example illustrates one usage of `ListObjects`.

```
<?xml version="1.0" encoding="UTF-8"?>
<ListBucketResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Name>bucket</Name>
  <Prefix/>
  <Marker/>
  <MaxKeys>1000</MaxKeys>
  <IsTruncated>>false</IsTruncated>
  <Contents>
    <Key>my-image.jpg</Key>
    <LastModified>2009-10-12T17:50:30.000Z</LastModified>
    <ETag>"fba9dede5f27731c9771645a39863328"</ETag>
    <Size>434234</Size>
    <StorageClass>STANDARD</StorageClass>
    <Owner>
      <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeebf76c078efc7c6caea54ba06a</ID>
    </Owner>
  </Contents>
  <Contents>
    <Key>my-third-image.jpg</Key>
    <LastModified>2009-10-12T17:50:30.000Z</LastModified>
    <ETag>"1b2cf535f27731c974343645a3985328"</ETag>
    <Size>64994</Size>
    <StorageClass>STANDARD_IA</StorageClass>
    <Owner>
      <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeebf76c078efc7c6caea54ba06a</ID>
    </Owner>
  </Contents>
</ListBucketResult>
```

Sample Request: Using request parameters

This example lists up to 40 keys in the quotes bucket that start with N and occur lexicographically after Ned.

```
GET /?prefix=N&marker=Ned&max-keys=40 HTTP/1.1
Host: quotes.s3.<Region>.amazonaws.com
Date: Wed, 01 Mar 2006 12:00:00 GMT
Authorization: authorization string
```

Sample Response

This example illustrates one usage of ListObjects.

```
HTTP/1.1 200 OK
x-amz-id-2: gyB+3jRPnrkN98ZajxHXr3u7EFM67bNgSAXexeEHndCX/7GRnfTXxReKUQF28IfP
x-amz-request-id: 3B3C7C725673C630
Date: Wed, 01 Mar 2006 12:00:00 GMT
Content-Type: application/xml
Content-Length: 302
Connection: close
Server: AmazonS3
```

```
<?xml version="1.0" encoding="UTF-8"?>
<ListBucketResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Name>quotes</Name>
  <Prefix>N</Prefix>
  <Marker>Ned</Marker>
  <MaxKeys>40</MaxKeys>
  <IsTruncated>false</IsTruncated>
  <Contents>
    <Key>Nelson</Key>
    <LastModified>2006-01-01T12:00:00.000Z</LastModified>
    <ETag>"828ef3fdfa96f00ad9f27c383fc9ac7f"</ETag>
    <Size>5</Size>
    <StorageClass>STANDARD</StorageClass>
    <Owner>
      <ID>bcaf161ca5fb16fd081034f</ID>
    </Owner>
  </Contents>
  <Contents>
    <Key>Neo</Key>
    <LastModified>2006-01-01T12:00:00.000Z</LastModified>
    <ETag>"828ef3fdfa96f00ad9f27c383fc9ac7f"</ETag>
    <Size>4</Size>
    <StorageClass>STANDARD</StorageClass>
    <Owner>
      <ID>bcaf1ffd86a5fb16fd081034f</ID>
    </Owner>
  </Contents>
</ListBucketResult>
```


Sample Request: Using a prefix and delimiter

For this example, we assume that you have the following keys in your bucket:

- sample.jpg
- photos/2006/January/sample.jpg
- photos/2006/February/sample2.jpg
- photos/2006/February/sample3.jpg
- photos/2006/February/sample4.jpg

The following GET request specifies the `delimiter` parameter with a value of `/`.

```
GET /?delimiter=/ HTTP/1.1
Host: example-bucket.s3.<Region>.amazonaws.com
Date: Wed, 01 Mar 2006 12:00:00 GMT
Authorization: authorization string
```

Sample Response

The key `sample.jpg` does not contain the delimiter character, and Amazon S3 returns it in the `Contents` element in the response. However, all of the other keys contain the delimiter character. Amazon S3 groups these keys and returns a single `CommonPrefixes` element with the `Prefix` value `photos/`, which is a substring from the beginning of these keys to the first occurrence of the specified delimiter.

```
<ListBucketResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Name>example-bucket</Name>
  <Prefix></Prefix>
  <Marker></Marker>
  <MaxKeys>1000</MaxKeys>
  <Delimiter></Delimiter>
  <IsTruncated>>false</IsTruncated>
```

```
<Contents>
  <Key>sample.jpg</Key>
  <LastModified>2011-02-26T01:56:20.000Z</LastModified>
  <ETag>"bf1d737a4d46a19f3bcd6905cc8b902"</ETag>
  <Size>142863</Size>
  <Owner>
    <ID>canonical-user-id</ID>
  </Owner>
  <StorageClass>STANDARD</StorageClass>
</Contents>
<CommonPrefixes>
  <Prefix>photos/</Prefix>
</CommonPrefixes>
</ListBucketResult>
```

Sample Request

The following GET request specifies the `delimiter` parameter with the value `/`, and the `prefix` parameter with the value `photos/2006/`.

```
GET /?prefix=photos/2006/&delimiter=/ HTTP/1.1
Host: example-bucket.s3.<Region>.amazonaws.com
Date: Wed, 01 Mar 2006 12:00:00 GMT
Authorization: authorization string
```

Sample Response

In response, Amazon S3 returns only the keys that start with the specified prefix. Amazon S3 uses the delimiter character to group keys that contain the same substring until the first occurrence of the delimiter character after the specified prefix. For each such key group, Amazon S3 returns one `CommonPrefixes` element in the response. The keys grouped under this `CommonPrefixes` element are not returned elsewhere in the response. The value returned in the `CommonPrefixes` element is a substring that starts at the beginning of the key and ends at the first occurrence of the specified delimiter after the prefix.

```
<ListBucketResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Name>example-bucket</Name>
  <Prefix>photos/2006/</Prefix>
  <Marker></Marker>
  <MaxKeys>1000</MaxKeys>
  <Delimiter></Delimiter>
  <IsTruncated>false</IsTruncated>

  <CommonPrefixes>
    <Prefix>photos/2006/February/</Prefix>
  </CommonPrefixes>
  <CommonPrefixes>
    <Prefix>photos/2006/January/</Prefix>
  </CommonPrefixes>
</ListBucketResult>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListObjectsV2

Service: Amazon S3

Returns some or all (up to 1,000) of the objects in a bucket with each request. You can use the request parameters as selection criteria to return a subset of the objects in a bucket. A 200 OK response can contain valid or invalid XML. Make sure to design your application to parse the contents of the response and handle it appropriately. For more information about listing objects, see [Listing object keys programmatically](#) in the *Amazon S3 User Guide*. To get a list of your buckets, see [ListBuckets](#).

Note

- **General purpose bucket** - For general purpose buckets, ListObjectsV2 doesn't return prefixes that are related only to in-progress multipart uploads.
- **Directory buckets** - For directory buckets, ListObjectsV2 response includes the prefixes that are related only to in-progress multipart uploads.
- **Directory buckets** - For directory buckets, you must make requests for this API operation to the Zonal endpoint. These endpoints support virtual-hosted-style requests in the format `https://amzn-s3-demo-bucket.s3express-zone-id.region-code.amazonaws.com/key-name`. Path-style requests are not supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

Permissions

- **General purpose bucket permissions** - To use this operation, you must have READ access to the bucket. You must have permission to perform the `s3:ListBucket` action. The bucket owner has this permission by default and can grant this permission to others. For more information about permissions, see [Permissions Related to Bucket Subresource Operations](#) and [Managing Access Permissions to Your Amazon S3 Resources](#) in the *Amazon S3 User Guide*.
- **Directory bucket permissions** - To grant access to this API operation on a directory bucket, we recommend that you use the [CreateSession](#) API operation for session-based authorization. Specifically, you grant the `s3express:CreateSession` permission to the directory bucket in a bucket policy or an IAM identity-based policy. Then, you make the

CreateSession API call on the bucket to obtain a session token. With the session token in your request header, you can make API requests to this operation. After the session token expires, you make another CreateSession API call to generate a new session token for use. AWS CLI or SDKs create session and refresh the session token automatically to avoid service interruptions when a session expires. For more information about authorization, see [CreateSession](#).

Sorting order of returned objects

- **General purpose bucket** - For general purpose buckets, ListObjectsV2 returns objects in lexicographical order based on their key names.
- **Directory bucket** - For directory buckets, ListObjectsV2 does not return objects in lexicographical order.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `Bucket-name.s3express-zone-id.region-code.amazonaws.com`.

Important

This section describes the latest revision of this action. We recommend that you use this revised API operation for application development. For backward compatibility, Amazon S3 continues to support the prior version of this API operation, [ListObjects](#).

The following operations are related to ListObjectsV2:

- [GetObject](#)
- [PutObject](#)
- [CreateBucket](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?list-type=2&continuation-token=ContinuationToken&delimiter=Delimiter&encoding-  
type=EncodingType&fetch-owner=FetchOwner&max-keys=MaxKeys&prefix=Prefix&start-  
after=StartAfter HTTP/1.1  
Host: Bucket.s3.amazonaws.com  
x-amz-request-payer: RequestPayer  
x-amz-expected-bucket-owner: ExpectedBucketOwner  
x-amz-optional-object-attributes: OptionalObjectAttributes
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

Directory buckets - When you use this operation with a directory bucket, you must use virtual-hosted-style requests in the format *Bucket-name.s3express-zone-id.region-code.amazonaws.com*. Path-style requests are not supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must follow the format *bucket-base-name--zone-id--x-s3* (for example, *amzn-s3-demo-bucket--usw2-az1--x-s3*). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form *AccessPointName-AccountId.s3-accesspoint.Region.amazonaws.com*. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

Note

Object Lambda access points are not supported by directory buckets.

S3 on Outposts - When you use this action with S3 on Outposts, you must direct requests to the S3 on Outposts hostname. The S3 on Outposts hostname takes the form `AccessPointName-AccountId.outpostID.s3-outposts.Region.amazonaws.com`. When you use this action with S3 on Outposts, the destination bucket must be the Outposts access point ARN or the access point alias. For more information about S3 on Outposts, see [What is S3 on Outposts?](#) in the *Amazon S3 User Guide*.

Required: Yes

continuation-token

ContinuationToken indicates to Amazon S3 that the list is being continued on this bucket with a token. ContinuationToken is obfuscated and is not a real key. You can use this ContinuationToken for pagination of the list results.

delimiter

A delimiter is a character that you use to group keys.

CommonPrefixes is filtered out from results if it is not lexicographically greater than the StartAfter value.

Note

- **Directory buckets** - For directory buckets, / is the only supported delimiter.
- **Directory buckets** - When you query ListObjectsV2 with a delimiter during in-progress multipart uploads, the CommonPrefixes response parameter contains the prefixes that are associated with the in-progress multipart uploads. For more information about multipart uploads, see [Multipart Upload Overview](#) in the *Amazon S3 User Guide*.

encoding-type

Encoding type used by Amazon S3 to encode the [object keys](#) in the response. Responses are encoded only in UTF-8. An object key can contain any Unicode character. However, the XML 1.0 parser can't parse certain characters, such as characters with an ASCII value from 0 to 10. For characters that aren't supported in XML 1.0, you can add this parameter to request that Amazon S3 encode the keys in the response. For more information about characters to avoid in object key names, see [Object key naming guidelines](#).

Note

When using the URL encoding type, non-ASCII characters that are used in an object's key name will be percent-encoded according to UTF-8 code values. For example, the object `test_file(3).png` will appear as `test_file%283%29.png`.

Valid Values: `url`

fetch-owner

The `owner` field is not present in `ListObjectsV2` by default. If you want to return the `owner` field with each key in the result, then set the `FetchOwner` field to `true`.

Note

Directory buckets - For directory buckets, the bucket owner is returned as the object owner for all objects.

max-keys

Sets the maximum number of keys returned in the response. By default, the action returns up to 1,000 key names. The response might contain fewer keys but will never contain more.

prefix

Limits the response to keys that begin with the specified prefix.

Note

Directory buckets - For directory buckets, only prefixes that end in a delimiter (`/`) are supported.

start-after

`StartAfter` is where you want Amazon S3 to start listing from. Amazon S3 starts listing after this specified key. `StartAfter` can be any key in the bucket.

Note

This functionality is not supported for directory buckets.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-optional-object-attributes

Specifies the optional fields that you want returned in the response. Fields that you do not specify are not returned.

Note

This functionality is not supported for directory buckets.

Valid Values: RestoreStatus

x-amz-request-payer

Confirms that the requester knows that she or he will be charged for the list objects request in V2 style. Bucket owners need not specify this parameter in their requests.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

Request Body

The request does not have a request body.

Response Syntax

```

HTTP/1.1 200
x-amz-request-charged: RequestCharged
<?xml version="1.0" encoding="UTF-8"?>
<ListBucketResult>
  <IsTruncated>boolean</IsTruncated>
  <Contents>
    <ChecksumAlgorithm>string</ChecksumAlgorithm>
    ...
    <ChecksumType>string</ChecksumType>
    <ETag>string</ETag>
    <Key>string</Key>
    <LastModified>timestamp</LastModified>
    <Owner>
      <DisplayName>string</DisplayName>
      <ID>string</ID>
    </Owner>
    <RestoreStatus>
      <IsRestoreInProgress>boolean</IsRestoreInProgress>
      <RestoreExpiryDate>timestamp</RestoreExpiryDate>
    </RestoreStatus>
    <Size>long</Size>
    <StorageClass>string</StorageClass>
  </Contents>
  ...
  <Name>string</Name>
  <Prefix>string</Prefix>
  <Delimiter>string</Delimiter>
  <MaxKeys>integer</MaxKeys>
  <CommonPrefixes>
    <Prefix>string</Prefix>
  </CommonPrefixes>
  ...
  <EncodingType>string</EncodingType>
  <KeyCount>integer</KeyCount>
  <ContinuationToken>string</ContinuationToken>
  <NextContinuationToken>string</NextContinuationToken>
  <StartAfter>string</StartAfter>
</ListBucketResult>

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: `requester`

The following data is returned in XML format by the service.

ListBucketResult

Root level tag for the ListBucketResult parameters.

Required: Yes

CommonPrefixes

All of the keys (up to 1,000) that share the same prefix are grouped together. When counting the total numbers of returns by this API operation, this group of keys is considered as one item.

A response can contain CommonPrefixes only if you specify a delimiter.

CommonPrefixes contains all (if there are any) keys between Prefix and the next occurrence of the string specified by a delimiter.

CommonPrefixes lists keys that act like subdirectories in the directory specified by Prefix.

For example, if the prefix is `notes/` and the delimiter is a slash (`/`) as in `notes/summer/july`, the common prefix is `notes/summer/`. All of the keys that roll up into a common prefix count as a single return when calculating the number of returns.

Note

- **Directory buckets** - For directory buckets, only prefixes that end in a delimiter (/) are supported.
- **Directory buckets** - When you query ListObjectsV2 with a delimiter during in-progress multipart uploads, the CommonPrefixes response parameter contains the prefixes that are associated with the in-progress multipart uploads. For more information about multipart uploads, see [Multipart Upload Overview](#) in the *Amazon S3 User Guide*.

Type: Array of [CommonPrefix](#) data types

Contents

Metadata about each object returned.

Type: Array of [Object](#) data types

ContinuationToken

If ContinuationToken was sent with the request, it is included in the response. You can use the returned ContinuationToken for pagination of the list response.

Type: String

Delimiter

Causes keys that contain the same string between the prefix and the first occurrence of the delimiter to be rolled up into a single result element in the CommonPrefixes collection. These rolled-up keys are not returned elsewhere in the response. Each rolled-up result counts as only one return against the MaxKeys value.

Note

Directory buckets - For directory buckets, / is the only supported delimiter.

Type: String

EncodingType

Encoding type used by Amazon S3 to encode object key names in the XML response.

If you specify the `encoding-type` request parameter, Amazon S3 includes this element in the response, and returns encoded key name values in the following response elements:

`Delimiter`, `Prefix`, `Key`, and `StartAfter`.

Type: String

Valid Values: `url`

IsTruncated

Set to `false` if all of the results were returned. Set to `true` if more keys are available to return. If the number of results exceeds that specified by `MaxKeys`, all of the results might not be returned.

Type: Boolean

KeyCount

`KeyCount` is the number of keys returned with this request. `KeyCount` will always be less than or equal to the `MaxKeys` field. For example, if you ask for 50 keys, your result will include 50 keys or fewer.

Type: Integer

MaxKeys

Sets the maximum number of keys returned in the response. By default, the action returns up to 1,000 key names. The response might contain fewer keys but will never contain more.

Type: Integer

Name

The bucket name.

Type: String

NextContinuationToken

`NextContinuationToken` is sent when `isTruncated` is `true`, which means there are more keys in the bucket that can be listed. The next list requests to Amazon S3 can be continued with this `NextContinuationToken`. `NextContinuationToken` is obfuscated and is not a real key

Type: String

Prefix

Keys that begin with the indicated prefix.

Note

Directory buckets - For directory buckets, only prefixes that end in a delimiter (/) are supported.

Type: String

StartAfter

If StartAfter was sent with the request, it is included in the response.

Note

This functionality is not supported for directory buckets.

Type: String

Errors

NoSuchBucket

The specified bucket does not exist.

HTTP Status Code: 404

Examples

Sample Request for general purpose buckets: Listing keys

This request returns the objects in bucket. The request specifies the `list-type` parameter, which indicates version 2 of the API operation.

```
GET /?list-type=2 HTTP/1.1
```

```
Host: bucket.s3.<Region>.amazonaws.com
x-amz-date: 20160430T233541Z
Authorization: authorization string
Content-Type: text/plain
```

Sample Response for general purpose buckets

This example illustrates one usage of ListObjectsV2.

```
<?xml version="1.0" encoding="UTF-8"?>
<ListBucketResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Name>bucket</Name>
  <Prefix/>
  <KeyCount>205</KeyCount>
  <MaxKeys>1000</MaxKeys>
  <IsTruncated>false</IsTruncated>
  <Contents>
    <Key>my-image.jpg</Key>
    <LastModified>2009-10-12T17:50:30.000Z</LastModified>
    <ETag>"fba9dede5f27731c9771645a39863328"</ETag>
    <Size>434234</Size>
    <StorageClass>STANDARD</StorageClass>
  </Contents>
  <Contents>
    ...
  </Contents>
  ...
</ListBucketResult>
```

Sample Request for general purpose buckets: Listing keys using the max-keys, prefix, and start-after parameters

In addition to the `list-type` parameter that indicates version 2 of the API operation, the request also specifies additional parameters to retrieve up to three keys in the quotes bucket that start with E and occur lexicographically after `ExampleGuide.pdf`.

```
GET /?list-type=2&max-keys=3&prefix=E&start-after=ExampleGuide.pdf HTTP/1.1
Host: quotes.s3.<Region>.amazonaws.com
```

```
x-amz-date: 20160430T232933Z
Authorization: authorization string
```

Sample Response for general purpose buckets

This example illustrates one usage of ListObjectsV2.

```
HTTP/1.1 200 OK
x-amz-id-2: gyB+3jRPnrkN98ZajxHXr3u7EFM67bNgSAXexeEHndCX/7GRnfTXxReKUQF28IfP
x-amz-request-id: 3B3C7C725673C630
Date: Sat, 30 Apr 2016 23:29:37 GMT
Content-Type: application/xml
Content-Length: length
Connection: close
Server: AmazonS3

<?xml version="1.0" encoding="UTF-8"?>
<ListBucketResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Name>quotes</Name>
  <Prefix>E</Prefix>
  <StartAfter>ExampleGuide.pdf</StartAfter>
  <KeyCount>1</KeyCount>
  <MaxKeys>3</MaxKeys>
  <IsTruncated>false</IsTruncated>
  <Contents>
    <Key>ExampleObject.txt</Key>
    <LastModified>2013-09-17T18:07:53.000Z</LastModified>
    <ETag>"599bab3ed2c697f1d26842727561fd94"</ETag>
    <Size>857</Size>
    <StorageClass>REDUCED_REDUNDANCY</StorageClass>
  </Contents>
</ListBucketResult>
```

Sample Request for general purpose buckets: Listing keys by using the prefix and delimiter parameters

This example illustrates the use of the `prefix` and the `delimiter` parameters in the request. For this example, we assume that you have the following keys in your bucket:

- sample.jpg
- photos/2006/January/sample.jpg
- photos/2006/February/sample2.jpg
- photos/2006/February/sample3.jpg
- photos/2006/February/sample4.jpg

The following GET request specifies the `delimiter` parameter with a value of `/`.

```
GET /?list-type=2&delimiter=/ HTTP/1.1
Host: example-bucket.s3.<Region>.amazonaws.com
x-amz-date: 20160430T235931Z
Authorization: authorization string
```

Sample Response for general purpose buckets

The key `sample.jpg` does not contain the delimiter character, and Amazon S3 returns it in the `Contents` element in the response. However, all of the other keys contain the delimiter character. Amazon S3 groups these keys and returns a single `CommonPrefixes` element with the `Prefix` value `photos/`. The `Prefix` element is a substring that starts at the beginning of these keys and ends at the first occurrence of the specified delimiter.

```
<ListBucketResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Name>example-bucket</Name>
  <Prefix></Prefix>
  <KeyCount>2</KeyCount>
  <MaxKeys>1000</MaxKeys>
  <Delimiter></Delimiter>
  <IsTruncated>false</IsTruncated>
  <Contents>
    <Key>sample.jpg</Key>
    <LastModified>2011-02-26T01:56:20.000Z</LastModified>
    <ETag>"bf1d737a4d46a19f3bcd6905cc8b902"</ETag>
    <Size>142863</Size>
    <StorageClass>STANDARD</StorageClass>
  </Contents>
  <CommonPrefixes>
    <Prefix>photos/</Prefix>
```

```
</CommonPrefixes>  
</ListBucketResult>
```

Sample Request for general purpose buckets

The following request specifies the `delimiter` parameter with the value `/`, and the `prefix` parameter with the value `photos/2006/`.

```
GET /?list-type=2&prefix=photos/2006/&delimiter=/ HTTP/1.1  
Host: example-bucket.s3.<Region>.amazonaws.com  
x-amz-date: 20160501T000433Z  
Authorization: authorization string
```

Sample Response for general purpose buckets

In response, Amazon S3 returns only the keys that start with the specified prefix. Further, Amazon S3 uses the delimiter character to group keys that contain the same substring until the first occurrence of the delimiter character after the specified prefix. For each such key group, Amazon S3 returns one `CommonPrefixes` element in the response. The keys grouped under this `CommonPrefixes` element are not returned elsewhere in the response. The `Prefix` value returned in the `CommonPrefixes` element is a substring that starts at the beginning of the key and ends at the first occurrence of the specified delimiter after the prefix.

Note

If you created folders by using the Amazon S3 console, you will see an additional 0-byte object with a key of `photos/2006/`. This object is created because of the way that the console supports folder structures. For more information, see [Organizing objects in the Amazon S3 console using folders](#) in the *Amazon S3 User Guide*.

```
<ListBucketResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">  
  <Name>example-bucket</Name>  
  <Prefix>photos/2006/</Prefix>  
  <KeyCount>2</KeyCount>  
  <MaxKeys>1000</MaxKeys>
```

```

<Delimiter>/</Delimiter>
<IsTruncated>>false</IsTruncated>

<CommonPrefixes>
  <Prefix>photos/2006/February/</Prefix>
</CommonPrefixes>
<CommonPrefixes>
  <Prefix>photos/2006/January/</Prefix>
</CommonPrefixes>
</ListBucketResult>

```

Sample Request for general purpose buckets: Using a continuation token

In this example, the initial request returns more than 1,000 keys. In response to this request, Amazon S3 returns the `IsTruncated` element with the value set to `true` and with a `NextContinuationToken` element.

```

GET /?list-type=2 HTTP/1.1
Host: bucket.s3.<Region>.amazonaws.com
Date: Mon, 02 May 2016 23:17:07 GMT
Authorization: authorization string

```

Sample Response for general purpose buckets: Using a continuation token

This example illustrates one usage of `ListObjectsV2`.

```

HTTP/1.1 200 OK
x-amz-id-2: gyB+3jRPnrkN98ZajxHXr3u7EFM67bNgSAXexeEHndCX/7GRnfTXxReKUQF28IfP
x-amz-request-id: 3B3C7C725673C630
Date: Sat, 30 Apr 2016 23:29:37 GMT
Content-Type: application/xml
Content-Length: length
Connection: close
Server: AmazonS3

<ListBucketResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Name>bucket</Name>

```

```

<Prefix></Prefix>
<NextContinuationToken>1ueGcxLPRx1Tr/XYExHnhbYLgveDs2J/wm36Hy4vb0wM=</
NextContinuationToken>
<KeyCount>1000</KeyCount>
<MaxKeys>1000</MaxKeys>
<IsTruncated>true</IsTruncated>
<Contents>
  <Key>happyface.jpg</Key>
  <LastModified>2014-11-21T19:40:05.000Z</LastModified>
  <ETag>"70ee1738b6b21e2c8a43f3a5ab0eee71"</ETag>
  <Size>11</Size>
  <StorageClass>STANDARD</StorageClass>
</Contents>
  ...
</ListBucketResult>

```

Sample request for general purpose buckets

In the following subsequent request, we include a continuation-token query parameter in the request with the value of the NextContinuationToken element from the preceding response.

```

GET /?list-type=2 HTTP/1.1
GET /?list-type=2&continuation-token=1ueGcxLPRx1Tr/XYExHnhbYLgveDs2J/wm36Hy4vb0wM=
HTTP/1.1

Host: bucket.s3.<Region>.amazonaws.com
Date: Mon, 02 May 2016 23:17:07 GMT
Authorization: authorization string

```

Sample response for general purpose buckets:

Amazon S3 returns a list of the next set of keys starting where the previous request ended.

```

HTTP/1.1 200 OK
x-amz-id-2: gyB+3jRPnrkN98ZajxHXr3u7EFM67bNgSAXexeEHndCX/7GRnfTXxReKUQF28IfP
x-amz-request-id: 3B3C7C725673C630

```

Date: Sat, 30 Apr 2016 23:29:37 GMT

Content-Type: application/xml

Content-Length: length

Connection: close

Server: AmazonS3

```
<ListBucketResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Name>bucket</Name>
  <Prefix></Prefix>
  <ContinuationToken>1ueGcxLPRx1Tr/XYExHnhbYLgveDs2J/wm36Hy4vb0wM=</ContinuationToken>
  <KeyCount>112</KeyCount>
  <MaxKeys>1000</MaxKeys>
  <IsTruncated>false</IsTruncated>
  <Contents>
    <Key>happyfacex.jpg</Key>
    <LastModified>2014-11-21T19:40:05.000Z</LastModified>
    <ETag>"70ee1738b6b21e2c8a43f3a5ab0eee71"</ETag>
    <Size>1111</Size>
    <StorageClass>STANDARD</StorageClass>
  </Contents>
  ...
</ListBucketResult>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListObjectVersions

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Returns metadata about all versions of the objects in a bucket. You can also use request parameters as selection criteria to return metadata about a subset of all the object versions.

Important

To use this operation, you must have permission to perform the `s3:ListBucketVersions` action. Be aware of the name difference.

Note

A 200 OK response can contain valid or invalid XML. Make sure to design your application to parse the contents of the response and handle it appropriately.

To use this operation, you must have READ access to the bucket.

The following operations are related to `ListObjectVersions`:

- [ListObjectsV2](#)
- [GetObject](#)
- [PutObject](#)
- [DeleteObject](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /?versions&delimiter=Delimiter&encoding-type=EncodingType&key-marker=KeyMarker&max-keys=MaxKeys&prefix=Prefix&version-id-marker=VersionIdMarker HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
x-amz-request-payer: RequestPayer
x-amz-optional-object-attributes: OptionalObjectAttributes
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name that contains the objects.

Required: Yes

delimiter

A delimiter is a character that you specify to group keys. All keys that contain the same string between the prefix and the first occurrence of the delimiter are grouped under a single result element in `CommonPrefixes`. These groups are counted as one result against the `max-keys` limitation. These keys are not returned elsewhere in the response.

`CommonPrefixes` is filtered out from results if it is not lexicographically greater than the key-marker.

encoding-type

Encoding type used by Amazon S3 to encode the [object keys](#) in the response. Responses are encoded only in UTF-8. An object key can contain any Unicode character. However, the XML 1.0 parser can't parse certain characters, such as characters with an ASCII value from 0 to 10. For characters that aren't supported in XML 1.0, you can add this parameter to request that Amazon S3 encode the keys in the response. For more information about characters to avoid in object key names, see [Object key naming guidelines](#).

Note

When using the URL encoding type, non-ASCII characters that are used in an object's key name will be percent-encoded according to UTF-8 code values. For example, the object `test_file(3).png` will appear as `test_file%283%29.png`.

Valid Values: `url`

key-marker

Specifies the key to start with when listing objects in a bucket.

max-keys

Sets the maximum number of keys returned in the response. By default, the action returns up to 1,000 key names. The response might contain fewer keys but will never contain more. If additional keys satisfy the search criteria, but were not returned because `max-keys` was exceeded, the response contains `<isTruncated>true</isTruncated>`. To return the additional keys, see `key-marker` and `version-id-marker`.

prefix

Use this parameter to select only those keys that begin with the specified prefix. You can use prefixes to separate a bucket into different groupings of keys. (You can think of using `prefix` to make groups in the same way that you'd use a folder in a file system.) You can use `prefix` with `delimiter` to roll up numerous objects into a single result under `CommonPrefixes`.

version-id-marker

Specifies the object version you want to start listing from.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code `403 Forbidden` (access denied).

x-amz-optional-object-attributes

Specifies the optional fields that you want returned in the response. Fields that you do not specify are not returned.

Valid Values: `RestoreStatus`

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
x-amz-request-charged: RequestCharged
<?xml version="1.0" encoding="UTF-8"?>
<ListVersionsResult>
  <IsTruncated>boolean</IsTruncated>
  <KeyMarker>string</KeyMarker>
  <VersionIdMarker>string</VersionIdMarker>
  <NextKeyMarker>string</NextKeyMarker>
  <NextVersionIdMarker>string</NextVersionIdMarker>
  <Version>
    <ChecksumAlgorithm>string</ChecksumAlgorithm>
    ...
    <ChecksumType>string</ChecksumType>
    <ETag>string</ETag>
    <IsLatest>boolean</IsLatest>
    <Key>string</Key>
    <LastModified>timestamp</LastModified>
    <Owner>
      <DisplayName>string</DisplayName>
      <ID>string</ID>
```

```

    </Owner>
    <RestoreStatus>
      <IsRestoreInProgress>boolean</IsRestoreInProgress>
      <RestoreExpiryDate>timestamp</RestoreExpiryDate>
    </RestoreStatus>
    <Size>long</Size>
    <StorageClass>string</StorageClass>
    <VersionId>string</VersionId>
  </Version>
  ...
  <DeleteMarker>
    <IsLatest>boolean</IsLatest>
    <Key>string</Key>
    <LastModified>timestamp</LastModified>
    <Owner>
      <DisplayName>string</DisplayName>
      <ID>string</ID>
    </Owner>
    <VersionId>string</VersionId>
  </DeleteMarker>
  ...
  <Name>string</Name>
  <Prefix>string</Prefix>
  <Delimiter>string</Delimiter>
  <MaxKeys>integer</MaxKeys>
  <CommonPrefixes>
    <Prefix>string</Prefix>
  </CommonPrefixes>
  ...
  <EncodingType>string</EncodingType>
</ListVersionsResult>

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

The following data is returned in XML format by the service.

ListVersionsResult

Root level tag for the ListVersionsResult parameters.

Required: Yes

CommonPrefixes

All of the keys rolled up into a common prefix count as a single return when calculating the number of returns.

Type: Array of [CommonPrefix](#) data types

DeleteMarker

Container for an object that is a delete marker. To learn more about delete markers, see [Working with delete markers](#).

Type: Array of [DeleteMarkerEntry](#) data types

Delimiter

The delimiter grouping the included keys. A delimiter is a character that you specify to group keys. All keys that contain the same string between the prefix and the first occurrence of the delimiter are grouped under a single result element in `CommonPrefixes`. These groups are counted as one result against the max-keys limitation. These keys are not returned elsewhere in the response.

Type: String

EncodingType

Encoding type used by Amazon S3 to encode object key names in the XML response.

If you specify the `encoding-type` request parameter, Amazon S3 includes this element in the response, and returns encoded key name values in the following response elements:

`KeyMarker`, `NextKeyMarker`, `Prefix`, `Key`, and `Delimiter`.

Type: String

Valid Values: `url`

IsTruncated

A flag that indicates whether Amazon S3 returned all of the results that satisfied the search criteria. If your results were truncated, you can make a follow-up paginated request by using the `NextKeyMarker` and `NextVersionIdMarker` response parameters as a starting place in another request to return the rest of the results.

Type: Boolean

KeyMarker

Marks the last key returned in a truncated response.

Type: String

MaxKeys

Specifies the maximum number of objects to return.

Type: Integer

Name

The bucket name.

Type: String

NextKeyMarker

When the number of responses exceeds the value of `MaxKeys`, `NextKeyMarker` specifies the first key not returned that satisfies the search criteria. Use this value for the key-marker request parameter in a subsequent request.

Type: String

NextVersionIdMarker

When the number of responses exceeds the value of `MaxKeys`, `NextVersionIdMarker` specifies the first object version not returned that satisfies the search criteria. Use this value for the `version-id-marker` request parameter in a subsequent request.

Type: String

Prefix

Selects objects that start with the value supplied by this parameter.

Type: String

Version

Container for version information.

Type: Array of [ObjectVersion](#) data types

VersionIdMarker

Marks the last version of the key returned in a truncated response.

Type: String

Examples

Sample Request

The following request returns all of the versions of all of the objects in the specified bucket.

```
GET /?versions HTTP/1.1
Host: BucketName.s3.<Region>.amazonaws.com
Date: Wed, 28 Oct 2009 22:32:00 +0000
Authorization: authorization string (see Authenticating Requests (AWS Signature Version 4))
```

Sample Response

This example illustrates one usage of `ListObjectVersions`.

```
<?xml version="1.0" encoding="UTF-8"?>

<ListVersionsResult xmlns="http://s3.amazonaws.com/doc/2006-03-01">
  <Name>bucket</Name>
  <Prefix>my</Prefix>
  <KeyMarker/>
  <VersionIdMarker/>
  <MaxKeys>5</MaxKeys>
  <IsTruncated>false</IsTruncated>
  <Version>
    <Key>my-image.jpg</Key>
    <VersionId>3/L4kqtJl40Nr8X8gdRQBpUMLUo</VersionId>
    <IsLatest>true</IsLatest>
    <LastModified>2009-10-12T17:50:30.000Z</LastModified>
    <ETag>"fba9dede5f27731c9771645a39863328"</ETag>
    <Size>434234</Size>
    <StorageClass>STANDARD</StorageClass>
    <Owner>
      <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeebf76c078efc7c6caea54ba06a</ID>
    </Owner>
  </Version>
  <DeleteMarker>
    <Key>my-second-image.jpg</Key>
    <VersionId>03jppff543dhffds434rfsFDN943fdsFkdmqnh892</VersionId>
    <IsLatest>true</IsLatest>
    <LastModified>2009-11-12T17:50:30.000Z</LastModified>
    <Owner>
      <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeebf76c078efc7c6caea54ba06a</ID>
    </Owner>
  </DeleteMarker>
  <Version>
    <Key>my-second-image.jpg</Key>
    <VersionId>QUpfdndhfd8438MNFNDN93jdnJFkdmqnh893</VersionId>
    <IsLatest>false</IsLatest>
    <LastModified>2009-10-10T17:50:30.000Z</LastModified>
    <ETag>"9b2cf535f27731c974343645a3985328"</ETag>
    <Size>166434</Size>
    <StorageClass>STANDARD</StorageClass>
    <Owner>
      <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeebf76c078efc7c6caea54ba06a</ID>
    </Owner>
  </Version>
  <DeleteMarker>
    <Key>my-third-image.jpg</Key>
```

```

    <VersionId>03jpff543dhffds434rfsFDN943fdsFkdmqnh892</VersionId>
    <IsLatest>true</IsLatest>
    <LastModified>2009-10-15T17:50:30.000Z</LastModified>
    <Owner>
      <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeebf76c078efc7c6caea54ba06a</ID>
    </Owner>
  </DeleteMarker>
  <Version>
    <Key>my-third-image.jpg</Key>
    <VersionId>UI0RUnfndfhnw89493jJFJ</VersionId>
    <IsLatest>false</IsLatest>
    <LastModified>2009-10-11T12:50:30.000Z</LastModified>
    <ETag>"772cf535f27731c974343645a3985328"</ETag>
    <Size>64</Size>
    <StorageClass>STANDARD</StorageClass>
    <Owner>
      <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeebf76c078efc7c6caea54ba06a</ID>
    </Owner>
  </Version>
</ListVersionsResult>

```

Sample Request

The following request returns objects in the order that they were stored, returning the most recently stored object first, starting with the value for key-marker.

```

GET /?versions&key-marker=key2 HTTP/1.1
Host: s3.amazonaws.com
Pragma: no-cache
Accept: image/gif, image/x-xbitmap, image/jpeg, image/pjpeg, */*
Date: Thu, 10 Dec 2009 22:46:32 +0000
Authorization: signatureValue

```

Sample Response

This example illustrates one usage of ListObjectVersions.

```

<?xml version="1.0" encoding="UTF-8"?>

```



```
<ListVersionsResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Name>mtp-versioning-fresh</Name>
  <Prefix/>
  <KeyMarker>key2</KeyMarker>
  <VersionIdMarker/>
  <MaxKeys>1000</MaxKeys>
  <IsTruncated>false</IsTruncated>
  <Version>
    <Key>key3</Key>
    <VersionId>I5VhmK6CDDdQ5Pwfe1gcHZWmHDpcv7gfmfc29UBxsKU.</VersionId>
    <IsLatest>true</IsLatest>
    <LastModified>2009-12-09T00:19:04.000Z</LastModified>
    <ETag>"396fefef536d5ce46c7537ecf978a360"</ETag>
    <Size>217</Size>
    <Owner>
      <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeebf76c078efc7c6caea54ba06a</ID>
    </Owner>
    <StorageClass>STANDARD</StorageClass>
  </Version>
  <DeleteMarker>
    <Key>sourcekey</Key>
    <VersionId>qDhprLU80sAlCFLu2DWgXAEDgKzWarn-HS_JU0TvYqs.</VersionId>
    <IsLatest>true</IsLatest>
    <LastModified>2009-12-10T16:38:11.000Z</LastModified>
    <Owner>
      <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeebf76c078efc7c6caea54ba06a</ID>
    </Owner>
  </DeleteMarker>
  <Version>
    <Key>sourcekey</Key>
    <VersionId>wxxQ7ezLaL5JN2Sislq66Syxxo0k7uHTUpb9qiiMxNg.</VersionId>
    <IsLatest>false</IsLatest>
    <LastModified>2009-12-10T16:37:44.000Z</LastModified>
    <ETag>"396fefef536d5ce46c7537ecf978a360"</ETag>
    <Size>217</Size>
    <Owner>
      <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeebf76c078efc7c6caea54ba06a</ID>
    </Owner>
    <StorageClass>STANDARD</StorageClass>
  </Version>
</ListVersionsResult>
```

Sample Request Using the prefix Parameter

This example returns objects whose keys begin with source.

```
GET /?versions&prefix=source HTTP/1.1
Host: bucket.s3.<Region>.amazonaws.com
Date: Wed, 28 Oct 2009 22:32:00 +0000
Authorization: authorization string
```

Sample Response

This example illustrates one usage of ListObjectVersions.

```
<?xml version="1.0" encoding="UTF-8"?>
<ListVersionsResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Name>mtp-versioning-fresh</Name>
  <Prefix>source</Prefix>
  <KeyMarker/>
  <VersionIdMarker/>
  <MaxKeys>1000</MaxKeys>
  <IsTruncated>>false</IsTruncated>
  <DeleteMarker>
    <Key>sourcekey</Key>
    <VersionId>qDhprLU80sAlCFLu2DWgXAEDgKzWarn-HS_JU0TvYqs.</VersionId>
    <IsLatest>>true</IsLatest>
    <LastModified>2009-12-10T16:38:11.000Z</LastModified>
    <Owner>
      <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeebf76c078efc7c6caea54ba06a</ID>
    </Owner>
  </DeleteMarker>
  <Version>
    <Key>sourcekey</Key>
    <VersionId>wxXQ7ezLaL5JN2Sislq66Syxxo0k7uHTUpb9qiiMxNg.</VersionId>
    <IsLatest>>false</IsLatest>
    <LastModified>2009-12-10T16:37:44.000Z</LastModified>
    <ETag>"396fefef536d5ce46c7537ecf978a360"</ETag>
    <Size>217</Size>
    <Owner>
      <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeebf76c078efc7c6caea54ba06a</ID>
```

```
</Owner>
<StorageClass>STANDARD</StorageClass>
</Version>
</ListVersionsResult>
```

Sample Request: Using the key-marker and version-id-marker Parameters

The following example returns objects starting at the specified key (key-marker) and version ID (version-id-marker).

```
GET /?versions&key-marker=key3&version-id-marker=t46ZenlYTZBnj HTTP/1.1
Host: bucket.s3.<Region>.amazonaws.com
Date: Wed, 28 Oct 2009 22:32:00 +0000
Authorization: signatureValue
```

Sample Response

This example illustrates one usage of ListObjectVersions.

```
<?xml version="1.0" encoding="UTF-8"?>
<ListVersionsResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Name>mtp-versioning-fresh</Name>
  <Prefix/>
  <KeyMarker>key3</KeyMarker>
  <VersionIdMarker>t46ZenlYTZBnj</VersionIdMarker>
  <MaxKeys>1000</MaxKeys>
  <IsTruncated>false</IsTruncated>
  <DeleteMarker>
    <Key>sourcekey</Key>
    <VersionId>qDhprLU80sAlCFLu2DWgXAEDgKzWarn-HS_JU0TvYqs.</VersionId>
    <IsLatest>true</IsLatest>
    <LastModified>2009-12-10T16:38:11.000Z</LastModified>
    <Owner>
      <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeebf76c078efc7c6caea54ba06a</ID>
    </Owner>
  </DeleteMarker>
  <Version>
    <Key>sourcekey</Key>
```

```

<VersionId>wxXQ7ezLaL5JN2Sislq66Syxxo0k7uHTUpb9qiiMxNg.</VersionId>
<IsLatest>false</IsLatest>
<LastModified>2009-12-10T16:37:44.000Z</LastModified>
<ETag>"396fefef536d5ce46c7537ecf978a360"</ETag>
<Size>217</Size>
<Owner>
  <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeebf76c078efc7c6caea54ba06a</ID>
</Owner>
<StorageClass>STANDARD</StorageClass>
</Version>
</ListVersionsResult>

```

Sample Request: Using the key-marker, version-id-marker, and max-keys Parameters

The following request returns up to three (the value of max-keys) objects starting with the key specified by key-marker and the version ID specified by version-id-marker.

```

GET /?versions&key-marker=key3&version-id-marker=t46Z0menlYTzBnj&max-keys=3
Host: bucket.s3.<Region>.amazonaws.com
Date: Wed, 28 Oct 2009 22:32:00 +0000
Authorization: authorization string

```

Sample Response

This example illustrates one usage of ListObjectVersions.

```

<?xml version="1.0" encoding="UTF-8"?>
<ListVersionsResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Name>mtp-versioning-fresh</Name>
  <Prefix/>
  <KeyMarker>key3</KeyMarker>
  <VersionIdMarker>null</VersionIdMarker>
  <NextKeyMarker>key3</NextKeyMarker>
  <NextVersionIdMarker>d-d309mfjFrUmoQ0DBsVqmcMV150I.</NextVersionIdMarker>
  <MaxKeys>3</MaxKeys>
  <IsTruncated>true</IsTruncated>
  <Version>
    <Key>key3</Key>
    <VersionId>8XECiENpj8pydEDJdd-_VRrvaGKAH0aGMNW7tg6UViI.</VersionId>
  
```

```

<IsLatest>false</IsLatest>
<LastModified>2009-12-09T00:18:23.000Z</LastModified>
<ETag>"396fefef536d5ce46c7537ecf978a360"</ETag>
<Size>217</Size>
<Owner>
  <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeebf76c078efc7c6caea54ba06a</ID>
</Owner>
<StorageClass>STANDARD</StorageClass>
</Version>
<Version>
  <Key>key3</Key>
  <VersionId>d-d309mfjFri40QYukDozqBt3UmoQ0DBsVqmcMV150I.</VersionId>
  <IsLatest>false</IsLatest>
  <LastModified>2009-12-09T00:18:08.000Z</LastModified>
  <ETag>"396fefef536d5ce46c7537ecf978a360"</ETag>
  <Size>217</Size>
  <Owner>
    <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeebf76c078efc7c6caea54ba06a</ID>
  </Owner>
  <StorageClass>STANDARD</StorageClass>
</Version>
</ListVersionsResult>

```

Sample Request: Using the delimiter and prefix Parameters

Assume you have the following keys in your bucket, example-bucket.

photos/2006/January/sample.jpg

photos/2006/February/sample.jpg

photos/2006/March/sample.jpg

videos/2006/March/sample.wmv

sample.jpg

The following GET versions request specifies the `delimiter` parameter with the value `/`.

```

GET /?versions&delimiter=/ HTTP/1.1
Host: example-bucket.s3.<Region>.amazonaws.com

```

```
Date: Wed, 02 Feb 2011 20:34:56 GMT
Authorization: authorization string
```

Sample Response

The list of keys from the specified bucket is shown in the following response.

The response returns the `sample.jpg` key in a `Version` element. However, because all the other keys contain the specified delimiter, a distinct substring, from the beginning of the key to the first occurrence of the delimiter, from each of these keys is returned in a `CommonPrefixes` element. The key substrings, `photos/` and `videos/`, in the `CommonPrefixes` element indicate that there are one or more keys with these key prefixes.

This is a useful scenario if you use key prefixes for your objects to create a logical folder-like structure. In this case, you can interpret the result as the folders `photos/` and `videos/` have one or more objects.

```
<ListVersionsResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Name>mvbucketwithversionon1</Name>
  <Prefix></Prefix>
  <KeyMarker></KeyMarker>
  <VersionIdMarker></VersionIdMarker>
  <MaxKeys>1000</MaxKeys>
  <Delimiter></Delimiter>
  <IsTruncated>>false</IsTruncated>

  <Version>
    <Key>Sample.jpg</Key>
    <VersionId>toxMzQlBsGyGCz1YuMWMp90cdXLzq0CH</VersionId>
    <IsLatest>true</IsLatest>
    <LastModified>2011-02-02T18:46:20.000Z</LastModified>
    <ETag>"3305f2cfc46c0f04559748bb039d69ae"</ETag>
    <Size>3191</Size>
    <Owner>
      <ID>852b113e7a2f25102679df27bb0ae12b3f85be6f290b936c4393484be31bebcc</ID>
      <DisplayName>display-name</DisplayName>
    </Owner>
    <StorageClass>STANDARD</StorageClass>
  </Version>
```

```
<CommonPrefixes>
  <Prefix>photos/</Prefix>
</CommonPrefixes>
<CommonPrefixes>
  <Prefix>videos/</Prefix>
</CommonPrefixes>
</ListVersionsResult>
```

Example

In addition to the `delimiter` parameter, you can filter results by adding a `prefix` parameter as shown in the following request.

```
GET /?versions&prefix=photos/2006/&delimiter=/ HTTP/1.1
Host: example-bucket.s3.<Region>.amazonaws.com
Date: Wed, 02 Feb 2011 19:34:02 GMT
Authorization: authorization string
```

Example

In this case, the response will include only object keys that start with the specified prefix. The value returned in the `CommonPrefixes` element is a substring from the beginning of the key to the first occurrence of the specified delimiter after the prefix.

Note

If you created folders by using the Amazon S3 console, you will see an additional 0-byte object with a key of `photos/2006/`. This object is created because of the way that the console supports folder structures. For more information, see [Organizing objects in the Amazon S3 console using folders](#) in the *Amazon S3 User Guide*.

```
<?xml version="1.0" encoding="UTF-8"?>
<ListVersionsResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Name>example-bucket</Name>
  <Prefix>photos/2006/</Prefix>
```

```
<KeyMarker></KeyMarker>
<VersionIdMarker></VersionIdMarker>
<MaxKeys>1000</MaxKeys>
<Delimiter></Delimiter>
<IsTruncated>>false</IsTruncated>
<CommonPrefixes>
  <Prefix>photos/2006/February/</Prefix>
</CommonPrefixes>
<CommonPrefixes>
  <Prefix>photos/2006/January/</Prefix>
</CommonPrefixes>
<CommonPrefixes>
  <Prefix>photos/2006/March/</Prefix>
</CommonPrefixes>
</ListVersionsResult>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListParts

Service: Amazon S3

Lists the parts that have been uploaded for a specific multipart upload.

To use this operation, you must provide the `upload ID` in the request. You obtain this `uploadID` by sending the initiate multipart upload request through [CreateMultipartUpload](#).

The `ListParts` request returns a maximum of 1,000 uploaded parts. The limit of 1,000 parts is also the default value. You can restrict the number of parts in a response by specifying the `max-parts` request parameter. If your multipart upload consists of more than 1,000 parts, the response returns an `IsTruncated` field with the value of `true`, and a `NextPartNumberMarker` element. To list remaining uploaded parts, in subsequent `ListParts` requests, include the `part-number-marker` query string parameter and set its value to the `NextPartNumberMarker` field value from the previous response.

For more information on multipart uploads, see [Uploading Objects Using Multipart Upload](#) in the *Amazon S3 User Guide*.

Note

Directory buckets - For directory buckets, you must make requests for this API operation to the Zonal endpoint. These endpoints support virtual-hosted-style requests in the format `https://amzn-s3-demo-bucket.s3express-zone-id.region-code.amazonaws.com/key-name`. Path-style requests are not supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

Permissions

- **General purpose bucket permissions** - For information about permissions required to use the multipart upload API, see [Multipart Upload and Permissions](#) in the *Amazon S3 User Guide*.

If the upload was created using server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS) or dual-layer server-side encryption with AWS KMS keys (DSSE-KMS), you must have permission to the `kms:Decrypt` action for the `ListParts` request to succeed.

- **Directory bucket permissions** - To grant access to this API operation on a directory bucket, we recommend that you use the [CreateSession](#) API operation for session-based authorization. Specifically, you grant the `s3express:CreateSession` permission to the directory bucket in a bucket policy or an IAM identity-based policy. Then, you make the `CreateSession` API call on the bucket to obtain a session token. With the session token in your request header, you can make API requests to this operation. After the session token expires, you make another `CreateSession` API call to generate a new session token for use. AWS CLI or SDKs create session and refresh the session token automatically to avoid service interruptions when a session expires. For more information about authorization, see [CreateSession](#).

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `Bucket-name.s3express-zone-id.region-code.amazonaws.com`.

The following operations are related to `ListParts`:

- [CreateMultipartUpload](#)
- [UploadPart](#)
- [CompleteMultipartUpload](#)
- [AbortMultipartUpload](#)
- [GetObjectAttributes](#)
- [ListMultipartUploads](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
GET /Key+?max-parts=MaxParts&part-number-marker=PartNumberMarker&uploadId=UploadId
HTTP/1.1
Host: Bucket.s3.amazonaws.com
```

```
x-amz-request-payer: RequestPayer  
x-amz-expected-bucket-owner: ExpectedBucketOwner  
x-amz-server-side-encryption-customer-algorithm: SSECustomerAlgorithm  
x-amz-server-side-encryption-customer-key: SSECustomerKey  
x-amz-server-side-encryption-customer-key-MD5: SSECustomerKeyMD5
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket to which the parts are being uploaded.

Directory buckets - When you use this operation with a directory bucket, you must use virtual-hosted-style requests in the format `Bucket-name.s3express-zone-id.region-code.amazonaws.com`. Path-style requests are not supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must follow the format `bucket-base-name--zone-id--x-s3` (for example, `amzn-s3-demo-bucket--usw2-az1--x-s3`). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form `AccessPointName-AccountId.s3-accesspoint.Region.amazonaws.com`. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

Note

Object Lambda access points are not supported by directory buckets.

S3 on Outposts - When you use this action with S3 on Outposts, you must direct requests to the S3 on Outposts hostname. The S3 on Outposts hostname takes the form

`AccessPointName-AccountId.outpostID.s3-outposts.Region.amazonaws.com`.
When you use this action with S3 on Outposts, the destination bucket must be the Outposts access point ARN or the access point alias. For more information about S3 on Outposts, see [What is S3 on Outposts?](#) in the *Amazon S3 User Guide*.

Required: Yes

Key

Object key for which the multipart upload was initiated.

Length Constraints: Minimum length of 1.

Required: Yes

max-parts

Sets the maximum number of parts to return.

part-number-marker

Specifies the part after which listing should begin. Only parts with higher part numbers will be listed.

uploadId

Upload ID identifying the multipart upload whose parts are being listed.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

x-amz-server-side-encryption-customer-algorithm

The server-side encryption (SSE) algorithm used to encrypt the object. This parameter is needed only when the object was created using a checksum algorithm. For more information, see [Protecting data using SSE-C keys](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

x-amz-server-side-encryption-customer-key

The server-side encryption (SSE) customer managed key. This parameter is needed only when the object was created using a checksum algorithm. For more information, see [Protecting data using SSE-C keys](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

x-amz-server-side-encryption-customer-key-MD5

The MD5 server-side encryption (SSE) customer managed key. This parameter is needed only when the object was created using a checksum algorithm. For more information, see [Protecting data using SSE-C keys](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
x-amz-abort-date: AbortDate
x-amz-abort-rule-id: AbortRuleId
x-amz-request-charged: RequestCharged
<?xml version="1.0" encoding="UTF-8"?>
<ListPartsResult>
  <Bucket>string</Bucket>
  <Key>string</Key>
  <UploadId>string</UploadId>
  <PartNumberMarker>integer</PartNumberMarker>
  <NextPartNumberMarker>integer</NextPartNumberMarker>
  <MaxParts>integer</MaxParts>
  <IsTruncated>boolean</IsTruncated>
  <Part>
    <ChecksumCRC32>string</ChecksumCRC32>
    <ChecksumCRC32C>string</ChecksumCRC32C>
    <ChecksumCRC64NVME>string</ChecksumCRC64NVME>
    <ChecksumMD5>string</ChecksumMD5>
    <ChecksumSHA1>string</ChecksumSHA1>
    <ChecksumSHA256>string</ChecksumSHA256>
    <ChecksumSHA512>string</ChecksumSHA512>
    <ChecksumXXHASH128>string</ChecksumXXHASH128>
    <ChecksumXXHASH3>string</ChecksumXXHASH3>
    <ChecksumXXHASH64>string</ChecksumXXHASH64>
    <ETag>string</ETag>
    <LastModified>timestamp</LastModified>
    <PartNumber>integer</PartNumber>
    <Size>long</Size>
  </Part>
  ...
  <Initiator>
    <DisplayName>string</DisplayName>
    <ID>string</ID>
  </Initiator>
  <Owner>
    <DisplayName>string</DisplayName>
    <ID>string</ID>
```

```
</Owner>
<StorageClass>string</StorageClass>
<ChecksumAlgorithm>string</ChecksumAlgorithm>
<ChecksumType>string</ChecksumType>
</ListPartsResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

x-amz-abort-date

If the bucket has a lifecycle rule configured with an action to abort incomplete multipart uploads and the prefix in the lifecycle rule matches the object name in the request, then the response includes this header indicating when the initiated multipart upload will become eligible for abort operation. For more information, see [Aborting Incomplete Multipart Uploads Using a Bucket Lifecycle Configuration](#).

The response will also include the `x-amz-abort-rule-id` header that will provide the ID of the lifecycle configuration rule that defines this action.

Note

This functionality is not supported for directory buckets.

x-amz-abort-rule-id

This header is returned along with the `x-amz-abort-date` header. It identifies applicable lifecycle configuration rule that defines the action to abort incomplete multipart uploads.

Note

This functionality is not supported for directory buckets.

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

The following data is returned in XML format by the service.

ListPartsResult

Root level tag for the ListPartsResult parameters.

Required: Yes

Bucket

The name of the bucket to which the multipart upload was initiated. Does not return the access point ARN or access point alias if used.

Type: String

ChecksumAlgorithm

The algorithm that was used to create a checksum of the object.

Type: String

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 | XXHASH64 | XXHASH3 | XXHASH128

ChecksumType

The checksum type, which determines how part-level checksums are combined to create an object-level checksum for multipart objects. You can use this header response to verify that the checksum type that is received is the same checksum type that was specified in

CreateMultipartUpload request. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

Type: String

Valid Values: COMPOSITE | FULL_OBJECT

Initiator

Container element that identifies who initiated the multipart upload. If the initiator is an AWS account, this element provides the same information as the Owner element. If the initiator is an IAM User, this element provides the user ARN.

Type: [Initiator](#) data type

IsTruncated

Indicates whether the returned list of parts is truncated. A true value indicates that the list was truncated. A list can be truncated if the number of parts exceeds the limit returned in the MaxParts element.

Type: Boolean

Key

Object key for which the multipart upload was initiated.

Type: String

Length Constraints: Minimum length of 1.

MaxParts

Maximum number of parts that were allowed in the response.

Type: Integer

NextPartNumberMarker

When a list is truncated, this element specifies the last part in the list, as well as the value to use for the part-number-marker request parameter in a subsequent request.

Type: Integer

Owner

Container element that identifies the object owner, after the object is created. If multipart upload is initiated by an IAM user, this element provides the parent account ID.

Note

Directory buckets - The bucket owner is returned as the object owner for all the parts.

Type: [Owner](#) data type

[Part](#)

Container for elements related to a particular part. A response can contain zero or more Part elements.

Type: Array of [Part](#) data types

[PartNumberMarker](#)

Specifies the part after which listing should begin. Only parts with higher part numbers will be listed.

Type: Integer

[StorageClass](#)

The class of storage used to store the uploaded object.

Note

Directory buckets - Directory buckets only support EXPRESS_ONEZONE (the S3 Express One Zone storage class) in Availability Zones and ONEZONE_IA (the S3 One Zone-Infrequent Access storage class) in Dedicated Local Zones.

Type: String

Valid Values: STANDARD | REDUCED_REDUNDANCY | STANDARD_IA | ONEZONE_IA | INTELLIGENT_TIERING | GLACIER | DEEP_ARCHIVE | OUTPOSTS | GLACIER_IR | SNOW | EXPRESS_ONEZONE | FSX_OPENZFS | FSX_ONTAP | AWS_BACKUP_WARM | AWS_BACKUP_LOW_COST_WARM

[UploadId](#)

Upload ID identifying the multipart upload whose parts are being listed.

Type: String

Examples

Sample Request for general purpose buckets

Assume you have uploaded parts with sequential part numbers starting with 1. The following List Parts request specifies `max-parts` and `part-number-marker` query parameters. The request lists the first two parts that follow part number 1, that is, you will get parts 2 and 3 in the response. If more parts exist, the result is a truncated result and therefore the response will return an `IsTruncated` element with the value `true`. The response will also return the `NextPartNumberMarker` element with the value 3, which should be used for the value of the `part-number-marker` request query string parameter in the next ListParts request.

```
GET /example-object?
uploadId=XXBsb2FkIElEIGZvciBlbHZpbmcncyVcdS1tb3ZpZS5tMnRzEEEwbG9hZA&max-parts=2&part-
number-marker=1 HTTP/1.1
Host: example-bucket.s3.<Region>.amazonaws.com
Date: Mon, 1 Nov 2010 20:34:56 GMT
Authorization: authorization string
```

Sample Response for general purpose buckets

This example illustrates one usage of ListParts.

```
HTTP/1.1 200 OK
x-amz-id-2: Uuag1LuByRx9e6j50nimru9p04ZVKnJ2Qz7/C1NPcfTWAtRPfTa0Fg==
x-amz-request-id: 656c76696e6727732072657175657374
Date: Mon, 1 Nov 2010 20:34:56 GMT
Content-Length: 985
Connection: keep-alive
Server: AmazonS3

<?xml version="1.0" encoding="UTF-8"?>
<ListPartsResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Bucket>example-bucket</Bucket>
  <Key>example-object</Key>
  <UploadId>XXBsb2FkIElEIGZvciBlbHZpbmcncyVcdS1tb3ZpZS5tMnRzEEEwbG9hZA</UploadId>
  <Initiator>
    <ID>arn:aws:iam::111122223333:user/some-user-11116a31-17b5-4fb7-9df5-
b288870f11xx</ID>
```

```
</Initiator>
<Owner>
  <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeebf76c078efc7c6caea54ba06a</ID>

</Owner>
<StorageClass>STANDARD</StorageClass>
<PartNumberMarker>1</PartNumberMarker>
<NextPartNumberMarker>3</NextPartNumberMarker>
<MaxParts>2</MaxParts>
<IsTruncated>true</IsTruncated>
<Part>
  <PartNumber>2</PartNumber>
  <LastModified>2010-11-10T20:48:34.000Z</LastModified>
  <ETag>"7778aef83f66abc1fa1e8477f296d394"</ETag>
  <Size>10485760</Size>
</Part>
<Part>
  <PartNumber>3</PartNumber>
  <LastModified>2010-11-10T20:48:33.000Z</LastModified>
  <ETag>"aaaa18db4cc2f85cedef654fccc4a4x8"</ETag>
  <Size>10485760</Size>
</Part>
</ListPartsResult>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)

- [AWS SDK for Ruby V3](#)

PutBucketAbac

Service: Amazon S3

Sets the attribute-based access control (ABAC) property of the general purpose bucket. You must have `s3:PutBucketABAC` permission to perform this action. When you enable ABAC, you can use tags for access control on your buckets. Additionally, when ABAC is enabled, you must use the [TagResource](#) and [UntagResource](#) actions to manage tags on your buckets. You can no longer use the [PutBucketTagging](#) and [DeleteBucketTagging](#) actions to tag your bucket. For more information, see [Enabling ABAC in general purpose buckets](#).

Request Syntax

```
PUT /?abac HTTP/1.1
Host: Bucket.s3.amazonaws.com
Content-MD5: ContentMD5
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<AbacStatus xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Status>string</Status>
</AbacStatus>
```

URI Request Parameters

The request uses the following URI parameters.

[Bucket](#)

The name of the general purpose bucket.

Required: Yes

[Content-MD5](#)

The MD5 hash of the PutBucketAbac request body.

For requests made using the AWS Command Line Interface (CLI) or AWS SDKs, this field is calculated automatically.

[x-amz-expected-bucket-owner](#)

The AWS account ID of the general purpose bucket's owner.

[x-amz-sdk-checksum-algorithm](#)

Indicates the algorithm that you want Amazon S3 to use to create the checksum. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 | XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in XML format.

[AbacStatus](#)

Root level tag for the AbacStatus parameters.

Required: Yes

[Status](#)

The ABAC status of the general purpose bucket.

Type: String

Valid Values: Enabled | Disabled

Required: No

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutBucketAccelerateConfiguration

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Sets the accelerate configuration of an existing bucket. Amazon S3 Transfer Acceleration is a bucket-level feature that enables you to perform faster data transfers to Amazon S3.

To use this operation, you must have permission to perform the `s3:PutAccelerateConfiguration` action. The bucket owner has this permission by default. The bucket owner can grant this permission to others. For more information about permissions, see [Permissions Related to Bucket Subresource Operations](#) and [Managing Access Permissions to Your Amazon S3 Resources](#).

The Transfer Acceleration state of a bucket can be set to one of the following two values:

- Enabled – Enables accelerated data transfers to the bucket.
- Suspended – Disables accelerated data transfers to the bucket.

The [GetBucketAccelerateConfiguration](#) action returns the transfer acceleration state of a bucket.

After setting the Transfer Acceleration state of a bucket to Enabled, it might take up to thirty minutes before the data transfer rates to the bucket increase.

The name of the bucket used for Transfer Acceleration must be DNS-compliant and must not contain periods (".").

For more information about transfer acceleration, see [Transfer Acceleration](#).

The following operations are related to PutBucketAccelerateConfiguration:

- [GetBucketAccelerateConfiguration](#)
- [CreateBucket](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /?accelerate HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
<?xml version="1.0" encoding="UTF-8"?>
<AccelerateConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Status>string</Status>
</AccelerateConfiguration>
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket for which the accelerate configuration is set.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-sdk-checksum-algorithm

Indicates the algorithm used to create the checksum for the request when you use the SDK. This header will not provide any additional functionality if you don't use the SDK. When you send this header, there must be a corresponding `x-amz-checksum` or `x-amz-trailer` header sent. Otherwise, Amazon S3 fails the request with the HTTP status code 400 Bad Request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

If you provide an individual checksum, Amazon S3 ignores any provided `ChecksumAlgorithm` parameter.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 |
XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in XML format.

AccelerateConfiguration

Root level tag for the `AccelerateConfiguration` parameters.

Required: Yes

Status

Specifies the transfer acceleration status of the bucket.

Type: String

Valid Values: Enabled | Suspended

Required: No

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Sample Request: Add transfer acceleration configuration to set acceleration status

The following is an example of a `PUT /?accelerate` request that enables transfer acceleration for the bucket named `examplebucket`.

```
PUT /?accelerate HTTP/1.1
Host: examplebucket.s3.<Region>.amazonaws.com
Date: Mon, 11 Apr 2016 12:00:00 GMT
Authorization: authorization string
Content-Type: text/plain
Content-Length: length

<AccelerateConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Status>Enabled</Status>
</AccelerateConfiguration>
```

Sample Response

This example illustrates one usage of PutBucketAccelerateConfiguration.

```
HTTP/1.1 200 OK
x-amz-id-2: YgIPIfBiKa2bj0KMg95r/0zo3emzU4dzsD4rcKCHQUAdQkf3ShJT00pXUueF6QKo
x-amz-request-id: 236A8905248E5A01
Date: Mon, 11 Apr 2016 12:00:00 GMT
Content-Length: 0
Server: AmazonS3
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)

- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutBucketAcl

Service: Amazon S3

Important

End of support notice: As of October 1, 2025, Amazon S3 has discontinued support for Email Grantee Access Control Lists (ACLs). If you attempt to use an Email Grantee ACL in a request after October 1, 2025, the request will receive an HTTP 405 (Method Not Allowed) error.

This change affects the following AWS Regions: US East (N. Virginia), US West (N. California), US West (Oregon), Asia Pacific (Singapore), Asia Pacific (Sydney), Asia Pacific (Tokyo), Europe (Ireland), and South America (São Paulo).

Note

This operation is not supported for directory buckets.

Sets the permissions on an existing bucket using access control lists (ACL). For more information, see [Using ACLs](#). To set the ACL of a bucket, you must have the `WRITE_ACP` permission.

You can use one of the following two ways to set a bucket's permissions:

- Specify the ACL in the request body
- Specify permissions using request headers

Note

You cannot specify access permission using both the body and the request headers.

Depending on your application needs, you may choose to set the ACL on a bucket using either the request body or the headers. For example, if you have an existing application that updates a bucket ACL using the request body, then you can continue to use that approach.

Important

If your bucket uses the bucket owner enforced setting for S3 Object Ownership, ACLs are disabled and no longer affect permissions. You must use policies to grant access to your bucket and the objects in it. Requests to set ACLs or update ACLs fail and return the `AccessControlListNotSupported` error code. Requests to read ACLs are still supported. For more information, see [Controlling object ownership](#) in the *Amazon S3 User Guide*.

Permissions

You can set access permissions by using one of the following methods:

- Specify a canned ACL with the `x-amz-acl` request header. Amazon S3 supports a set of predefined ACLs, known as *canned ACLs*. Each canned ACL has a predefined set of grantees and permissions. Specify the canned ACL name as the value of `x-amz-acl`. If you use this header, you cannot use other access control-specific headers in your request. For more information, see [Canned ACL](#).
- Specify access permissions explicitly with the `x-amz-grant-read`, `x-amz-grant-read-acp`, `x-amz-grant-write-acp`, and `x-amz-grant-full-control` headers. When using these headers, you specify explicit access permissions and grantees (AWS accounts or Amazon S3 groups) who will receive the permission. If you use these ACL-specific headers, you cannot use the `x-amz-acl` header to set a canned ACL. These parameters map to the set of permissions that Amazon S3 supports in an ACL. For more information, see [Access Control List \(ACL\) Overview](#).

You specify each grantee as a `type=value` pair, where the `type` is one of the following:

- `id` – if the value specified is the canonical user ID of an AWS account
- `uri` – if you are granting permissions to a predefined group
- `emailAddress` – if the value specified is the email address of an AWS account

Note

Using email addresses to specify a grantee is only supported in the following AWS Regions:

- US East (N. Virginia)

- US West (N. California)
- US West (Oregon)
- Asia Pacific (Singapore)
- Asia Pacific (Sydney)
- Asia Pacific (Tokyo)
- Europe (Ireland)
- South America (São Paulo)

For a list of all the Amazon S3 supported Regions and endpoints, see [Regions and Endpoints](#) in the AWS General Reference.

For example, the following `x-amz-grant-write` header grants create, overwrite, and delete objects permission to LogDelivery group predefined by Amazon S3 and two AWS accounts identified by their email addresses.

```
x-amz-grant-write: uri="http://acs.amazonaws.com/groups/s3/LogDelivery", id="111122223333", id="555566667777"
```

You can use either a canned ACL or specify access permissions explicitly. You cannot do both.

Grantee Values

You can specify the person (grantee) to whom you're assigning access rights (using request elements) in the following ways. For examples of how to specify these grantee values in JSON format, see the AWS CLI example in [Enabling Amazon S3 server access logging](#) in the *Amazon S3 User Guide*.

- By the person's ID:

```
<Grantee xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:type="CanonicalUser"><ID><>ID<></ID><DisplayName><>GranteesEmail<></DisplayName> </Grantee>
```

DisplayName is optional and ignored in the request

- By URI:

```
<Grantee xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:type="Group"><URI><>http://acs.amazonaws.com/groups/global/AuthenticatedUsers<></URI></Grantee>
```


- By Email address:

```
<Grantee xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:type="AmazonCustomerByEmail"><EmailAddress><>Grantees@email.com<></
EmailAddress>&</Grantee>
```

The grantee is resolved to the CanonicalUser and, in a response to a GET Object acl request, appears as the CanonicalUser.

 Note

Using email addresses to specify a grantee is only supported in the following AWS Regions:

- US East (N. Virginia)
- US West (N. California)
- US West (Oregon)
- Asia Pacific (Singapore)
- Asia Pacific (Sydney)
- Asia Pacific (Tokyo)
- Europe (Ireland)
- South America (São Paulo)

For a list of all the Amazon S3 supported Regions and endpoints, see [Regions and Endpoints](#) in the AWS General Reference.

The following operations are related to PutBucketAcl:

- [CreateBucket](#)
- [DeleteBucket](#)
- [GetObjectAcl](#)

⚠ Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /?acl HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-acl: ACL
Content-MD5: ContentMD5
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-grant-full-control: GrantFullControl
x-amz-grant-read: GrantRead
x-amz-grant-read-acp: GrantReadACP
x-amz-grant-write: GrantWrite
x-amz-grant-write-acp: GrantWriteACP
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<AccessControlPolicy xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <AccessControlList>
    <Grant>
      <Grantee>
        <DisplayName>string</DisplayName>
        <EmailAddress>string</EmailAddress>
        <ID>string</ID>
        <xsi:type>string</xsi:type>
        <URI>string</URI>
      </Grantee>
      <Permission>string</Permission>
    </Grant>
  </AccessControlList>
  <Owner>
    <DisplayName>string</DisplayName>
    <ID>string</ID>
  </Owner>
</AccessControlPolicy>
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket to which to apply the ACL.

Required: Yes

Content-MD5

The Base64 encoded 128-bit MD5 digest of the data. This header must be used as a message integrity check to verify that the request body was not corrupted in transit. For more information, go to [RFC 1864](#).

For requests made using the AWS Command Line Interface (CLI) or AWS SDKs, this field is calculated automatically.

x-amz-acl

The canned ACL to apply to the bucket.

Valid Values: `private` | `public-read` | `public-read-write` | `authenticated-read`

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-grant-full-control

Allows grantee the read, write, read ACP, and write ACP permissions on the bucket.

x-amz-grant-read

Allows grantee to list the objects in the bucket.

x-amz-grant-read-acp

Allows grantee to read the bucket ACL.

x-amz-grant-write

Allows grantee to create new objects in the bucket.

For the bucket and object owners of existing objects, also allows deletions and overwrites of those objects.

[x-amz-grant-write-acp](#)

Allows grantee to write the ACL for the applicable bucket.

[x-amz-sdk-checksum-algorithm](#)

Indicates the algorithm used to create the checksum for the request when you use the SDK. This header will not provide any additional functionality if you don't use the SDK. When you send this header, there must be a corresponding `x-amz-checksum` or `x-amz-trailer` header sent. Otherwise, Amazon S3 fails the request with the HTTP status code `400 Bad Request`. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

If you provide an individual checksum, Amazon S3 ignores any provided `ChecksumAlgorithm` parameter.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 |
XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in XML format.

[AccessControlPolicy](#)

Root level tag for the `AccessControlPolicy` parameters.

Required: Yes

[Grants](#)

A list of grants.

Type: Array of [Grant](#) data types

Required: No

[Owner](#)

Container for the bucket owner's display name and ID.

Type: [Owner](#) data type

Required: No

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Sample Request: Access permissions specified in the body

The following request grants access permission to the existing `examplebucket` bucket. The request specifies the ACL in the body. In addition to granting full control to the bucket owner, the XML specifies the following grants.

- Grant the `AllUsers` group `READ` permission on the bucket.
- Grant the `LogDelivery` group `WRITE` permission on the bucket.
- Grant an AWS account, identified by email address, `WRITE_ACP` permission.
- Grant an AWS account, identified by canonical user ID, `READ_ACP` permission.

```
PUT ?acl HTTP/1.1
Host: examplebucket.s3.<Region>.amazonaws.com
Content-Length: 1660
x-amz-date: Thu, 12 Apr 2012 20:04:21 GMT
Authorization: authorization string

<AccessControlPolicy xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Owner>
    <ID>852b113e7a2f25102679df27bb0ae12b3f85be6BucketOwnerCanonicalUserID</ID>
    <DisplayName>OwnerDisplayName</DisplayName>
  </Owner>
  <AccessControlList>
    <Grant>
      <Grantee xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:type="CanonicalUser">
```

```

    <ID>852b113e7a2f25102679df27bb0ae12b3f85be6BucketOwnerCanonicalUserID</ID>
    <DisplayName>OwnerDisplayName</DisplayName>
  </Grantee>
  <Permission>FULL_CONTROL</Permission>
</Grant>
<Grant>
  <Grantee xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:type="Group">
    <URI xmlns="">http://acs.amazonaws.com/groups/global/AllUsers</URI>
  </Grantee>
  <Permission xmlns="">READ</Permission>
</Grant>
<Grant>
  <Grantee xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:type="Group">
    <URI xmlns="">http://acs.amazonaws.com/groups/s3/LogDelivery</URI>
  </Grantee>
  <Permission xmlns="">WRITE</Permission>
</Grant>
<Grant>
  <Grantee xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:type="AmazonCustomerByEmail">
    <EmailAddress xmlns="">xyz@amazon.com</EmailAddress>
  </Grantee>
  <Permission xmlns="">WRITE_ACP</Permission>
</Grant>
<Grant>
  <Grantee xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:type="CanonicalUser">
    <ID
xmlns="">f30716ab7115dcb44a5ef76e9d74b8e20567f63TestAccountCanonicalUserID</ID>
  </Grantee>
  <Permission xmlns="">READ_ACP</Permission>
</Grant>
</AccessControlList>
</AccessControlPolicy>

```

Sample Response

This example illustrates one usage of PutBucketAcl.

```

HTTP/1.1 200 OK
x-amz-id-2: Nxq03PNiMHXXGwjgv15LLgUoAmPvmG0xtZw2sxePXLhpIvcyouXDrcQUaWWXcOK0
x-amz-request-id: C651BC9B4E1BD401

```

```
Date: Thu, 12 Apr 2012 20:04:28 GMT
Content-Length: 0
Server: AmazonS3
```

Sample Request: Access permissions specified using headers

The following request uses ACL-specific request headers to grant the following permissions:

- Write permission to the Amazon S3 LogDelivery group and an AWS account identified by the email xyz@amazon.com.
- Read permission to the Amazon S3 AllUsers group

```
PUT ?acl HTTP/1.1
Host: examplebucket.s3.<Region>.amazonaws.com
x-amz-date: Sun, 29 Apr 2012 22:00:57 GMT
x-amz-grant-write: uri="http://acs.amazonaws.com/groups/s3/LogDelivery",
  emailAddress="xyz@amazon.com"
x-amz-grant-read: uri="http://acs.amazonaws.com/groups/global/AllUsers"
Accept: */*
Authorization: authorization string
```

Sample Response

This example illustrates one usage of PutBucketAcl.

```
HTTP/1.1 200 OK
x-amz-id-2: 0w9iImt23VF9s6Qof0TDzelF7mrryz7d04Mw23FQCi40205Zw28Zn+d340/RytoQ
x-amz-request-id: A6A8F01A38EC7138
Date: Sun, 29 Apr 2012 22:01:10 GMT
Content-Length: 0
Server: AmazonS3
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutBucketAnalyticsConfiguration

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Sets an analytics configuration for the bucket (specified by the analytics configuration ID). You can have up to 1,000 analytics configurations per bucket.

You can choose to have storage class analysis export analysis reports sent to a comma-separated values (CSV) flat file. See the `DataExport` request element. Reports are updated daily and are based on the object filters that you configure. When selecting data export, you specify a destination bucket and an optional destination prefix where the file is written. You can export the data to a destination bucket in a different account. However, the destination bucket must be in the same Region as the bucket that you are making the PUT analytics configuration to. For more information, see [Amazon S3 Analytics – Storage Class Analysis](#).

Important

You must create a bucket policy on the destination bucket where the exported file is written to grant permissions to Amazon S3 to write objects to the bucket. For an example policy, see [Granting Permissions for Amazon S3 Inventory and Storage Class Analysis](#).

To use this operation, you must have permissions to perform the `s3:PutAnalyticsConfiguration` action. The bucket owner has this permission by default. The bucket owner can grant this permission to others. For more information about permissions, see [Permissions Related to Bucket Subresource Operations](#) and [Managing Access Permissions to Your Amazon S3 Resources](#).

`PutBucketAnalyticsConfiguration` has the following special errors:

- *HTTP Error: HTTP 400 Bad Request*
 - *Code: InvalidArgument*
 - *Cause: Invalid argument.*
- *HTTP Error: HTTP 400 Bad Request*

- *Code: TooManyConfigurations*
- *Cause: You are attempting to create a new configuration but have already reached the 1,000-configuration limit.*
- *HTTP Error: HTTP 403 Forbidden*
- *Code: AccessDenied*
- *Cause: You are not the owner of the specified bucket, or you do not have the `s3:PutAnalyticsConfiguration` bucket permission to set the configuration on the bucket.*

The following operations are related to `PutBucketAnalyticsConfiguration`:

- [GetBucketAnalyticsConfiguration](#)
- [DeleteBucketAnalyticsConfiguration](#)
- [ListBucketAnalyticsConfigurations](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /?analytics&id=Id HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<AnalyticsConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Id>string</Id>
  <Filter>
    <And>
      <Prefix>string</Prefix>
      <Tag>
        <Key>string</Key>
        <Value>string</Value>
      </Tag>
    ...
```

```

</And>
<Prefix>string</Prefix>
<Tag>
  <Key>string</Key>
  <Value>string</Value>
</Tag>
</Filter>
<StorageClassAnalysis>
  <DataExport>
    <Destination>
      <S3BucketDestination>
        <Bucket>string</Bucket>
        <BucketAccountId>string</BucketAccountId>
        <Format>string</Format>
        <Prefix>string</Prefix>
      </S3BucketDestination>
    </Destination>
    <OutputSchemaVersion>string</OutputSchemaVersion>
  </DataExport>
</StorageClassAnalysis>
</AnalyticsConfiguration>

```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket to which an analytics configuration is stored.

Required: Yes

id

The ID that identifies the analytics configuration.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request accepts the following data in XML format.

[AnalyticsConfiguration](#)

Root level tag for the AnalyticsConfiguration parameters.

Required: Yes

[Filter](#)

The filter used to describe a set of objects for analyses. A filter must have exactly one prefix, one tag, or one conjunction (AnalyticsAndOperator). If no filter is provided, all objects will be considered in any analysis.

Type: [AnalyticsFilter](#) data type

Required: No

[Id](#)

The ID that identifies the analytics configuration.

Type: String

Required: Yes

[StorageClassAnalysis](#)

Contains data related to access patterns to be collected and made available to analyze the tradeoffs between different storage classes.

Type: [StorageClassAnalysis](#) data type

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Example 1: Creating an analytics configuration

The following PUT request for the bucket `examplebucket` creates a new or replaces an existing analytics configuration with the ID `report1`. The configuration is defined in the request body.

```
PUT /?analytics&id=report1 HTTP/1.1
Host: examplebucket.s3.<Region>.amazonaws.com
Date: Mon, 31 Oct 2016 12:00:00 GMT
Authorization: authorization string
Content-Length: length

<?xml version="1.0" encoding="UTF-8"?>
<AnalyticsConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Id>report1</Id>
  <Filter>
    <And>
      <Prefix>images/</Prefix>
      <Tag>
        <Key>dog</Key>
        <Value>corgi</Value>
      </Tag>
    </And>
  </Filter>
  <StorageClassAnalysis>
    <DataExport>
      <OutputSchemaVersion>V_1</OutputSchemaVersion>
      <Destination>
        <S3BucketDestination>
          <Format>CSV</Format>
          <BucketAccountId>123456789012</BucketAccountId>
          <Bucket>arn:aws:s3:::destination-bucket</Bucket>
          <Prefix>destination-prefix</Prefix>
        </S3BucketDestination>
      </Destination>
    </DataExport>
  </StorageClassAnalysis>
</AnalyticsConfiguration>
```

Sample Response

This example illustrates one usage of PutBucketAnalyticsConfiguration.

```
HTTP/1.1 200 OK
x-amz-id-2: YgIPIfBiKa2bj0KMg95r/0zo3emzU4dzsD4rcKCHQUAdQkf3ShJT00pXUueF6QKo
x-amz-request-id: 236A8905248E5A01
Date: Mon, 31 Oct 2016 12:00:00 GMT
Content-Length: 0
Server: AmazonS3
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutBucketCors

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Sets the `cors` configuration for your bucket. If the configuration exists, Amazon S3 replaces it.

To use this operation, you must be allowed to perform the `s3:PutBucketCORS` action. By default, the bucket owner has this permission and can grant it to others.

You set this configuration on a bucket so that the bucket can service cross-origin requests. For example, you might want to enable a request whose origin is `http://www.example.com` to access your Amazon S3 bucket at `my.example.bucket.com` by using the browser's XMLHttpRequest capability.

To enable cross-origin resource sharing (CORS) on a bucket, you add the `cors` subresource to the bucket. The `cors` subresource is an XML document in which you configure rules that identify origins and the HTTP methods that can be executed on your bucket. The document is limited to 64 KB in size.

When Amazon S3 receives a cross-origin request (or a pre-flight OPTIONS request) against a bucket, it evaluates the `cors` configuration on the bucket and uses the first `CORSRule` rule that matches the incoming browser request to enable a cross-origin request. For a rule to match, the following conditions must be met:

- The request's `Origin` header must match `AllowedOrigin` elements.
- The request method (for example, GET, PUT, HEAD, and so on) or the `Access-Control-Request-Method` header in case of a pre-flight OPTIONS request must be one of the `AllowedMethod` elements.
- Every header specified in the `Access-Control-Request-Headers` request header of a pre-flight request must match an `AllowedHeader` element.

For more information about CORS, go to [Enabling Cross-Origin Resource Sharing](#) in the *Amazon S3 User Guide*.

The following operations are related to PutBucketCors:

- [GetBucketCors](#)
- [DeleteBucketCors](#)
- [RESTOPTIONSobject](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /?cors HTTP/1.1
Host: Bucket.s3.amazonaws.com
Content-MD5: ContentMD5
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<CORSConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <CORSRule>
    <AllowedHeader>string</AllowedHeader>
    ...
    <AllowedMethod>string</AllowedMethod>
    ...
    <AllowedOrigin>string</AllowedOrigin>
    ...
    <ExposeHeader>string</ExposeHeader>
    ...
    <ID>string</ID>
    <MaxAgeSeconds>integer</MaxAgeSeconds>
  </CORSRule>
  ...
</CORSConfiguration>
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

Specifies the bucket impacted by the corsconfiguration.

Required: Yes

Content-MD5

The Base64 encoded 128-bit MD5 digest of the data. This header must be used as a message integrity check to verify that the request body was not corrupted in transit. For more information, go to [RFC 1864](#).

For requests made using the AWS Command Line Interface (CLI) or AWS SDKs, this field is calculated automatically.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-sdk-checksum-algorithm

Indicates the algorithm used to create the checksum for the request when you use the SDK. This header will not provide any additional functionality if you don't use the SDK. When you send this header, there must be a corresponding `x-amz-checksum` or `x-amz-trailer` header sent. Otherwise, Amazon S3 fails the request with the HTTP status code 400 Bad Request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

If you provide an individual checksum, Amazon S3 ignores any provided `ChecksumAlgorithm` parameter.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 | XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in XML format.

CORSConfiguration

Root level tag for the CORSConfiguration parameters.

Required: Yes

[CORSRule](#)

A set of origins and methods (cross-origin access that you want to allow). You can add up to 100 rules to the configuration.

Type: Array of [CORSRule](#) data types

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Example: CORS configuration on a bucket with two rules

- The first CORSRule allows cross-origin PUT, POST, and DELETE requests whose origin is `http://www.example.com` origins. The rule also allows all headers in a pre-flight OPTIONS request through the `Access-Control-Request-Headers` header. Therefore, in response to any pre-flight OPTIONS request, Amazon S3 will return any requested headers.
- The second rule allows cross-origin GET requests from all the origins. The '*' wildcard character refers to all origins.

```
<CORSConfiguration>
  <CORSRule>
    <AllowedOrigin>http://www.example.com</AllowedOrigin>

    <AllowedMethod>PUT</AllowedMethod>
    <AllowedMethod>POST</AllowedMethod>
    <AllowedMethod>DELETE</AllowedMethod>

    <AllowedHeader>*</AllowedHeader>
  </CORSRule>
```

```
<CORSRule>
  <AllowedOrigin>*</AllowedOrigin>
  <AllowedMethod>GET</AllowedMethod>
</CORSRule>
</CORSConfiguration>
```

Example: CORS configuration allows cross-origin PUT and POST requests from <http://www.example.com>

The cors configuration also allows additional optional configuration parameters as shown in the following cors configuration on a bucket. For example,

In the preceding configuration, CORSRule includes the following additional optional parameters:

- **MaxAgeSeconds**—Specifies the time in seconds that the browser will cache an Amazon S3 response to a pre-flight OPTIONS request for the specified resource. In this example, this parameter is 3000 seconds. Caching enables the browsers to avoid sending pre-flight OPTIONS request to Amazon S3 for repeated requests.
- **ExposeHeader**—Identifies the response header (in this case `x-amz-server-side-encryption`) that you want customers to be able to access from their applications (for example, from a JavaScript XMLHttpRequest object).

```
<CORSConfiguration>
  <CORSRule>
    <AllowedOrigin>http://www.example.com</AllowedOrigin>
    <AllowedMethod>PUT</AllowedMethod>
    <AllowedMethod>POST</AllowedMethod>
    <AllowedMethod>DELETE</AllowedMethod>
    <AllowedHeader>*</AllowedHeader>
    <MaxAgeSeconds>3000</MaxAgeSeconds>
    <ExposeHeader>x-amz-server-side-encryption</ExposeHeader>
  </CORSRule>
</CORSConfiguration>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)

- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutBucketEncryption

Service: Amazon S3

This operation configures default encryption and Amazon S3 Bucket Keys for an existing bucket. You can also [block encryption types](#) using this operation.

Note

Directory buckets - For directory buckets, you must make requests for this API operation to the Regional endpoint. These endpoints support path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name` . Virtual-hosted-style requests aren't supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

By default, all buckets have a default encryption configuration that uses server-side encryption with Amazon S3 managed keys (SSE-S3).

Note

- **General purpose buckets**
 - You can optionally configure default encryption for a bucket by using server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS) or dual-layer server-side encryption with AWS KMS keys (DSSE-KMS). If you specify default encryption by using SSE-KMS, you can also configure [Amazon S3 Bucket Keys](#). For information about the bucket default encryption feature, see [Amazon S3 Bucket Default Encryption](#) in the *Amazon S3 User Guide*.
 - If you use PutBucketEncryption to set your [default bucket encryption](#) to SSE-KMS, you should verify that your KMS key ID is correct. Amazon S3 doesn't validate the KMS key ID provided in PutBucketEncryption requests.
- **Directory buckets** - You can optionally configure default encryption for a bucket by using server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS).

- We recommend that the bucket's default encryption uses the desired encryption configuration and you don't override the bucket default encryption in your `CreateSession` requests or `PUT` object requests. Then, new objects are automatically encrypted with the desired encryption settings. For more information about the encryption overriding behaviors in directory buckets, see [Specifying server-side encryption with AWS KMS for new object uploads](#).
- Your SSE-KMS configuration can only support 1 [customer managed key](#) per directory bucket's lifetime. The [AWS managed key](#) (`aws/s3`) isn't supported.
- S3 Bucket Keys are always enabled for `GET` and `PUT` operations in a directory bucket and can't be disabled. S3 Bucket Keys aren't supported, when you copy SSE-KMS encrypted objects from general purpose buckets to directory buckets, from directory buckets to general purpose buckets, or between directory buckets, through [CopyObject](#), [UploadPartCopy](#), [the Copy operation in Batch Operations](#), or [the import jobs](#). In this case, Amazon S3 makes a call to AWS KMS every time a copy request is made for a KMS-encrypted object.
- When you specify an [AWS KMS customer managed key](#) for encryption in your directory bucket, only use the key ID or key ARN. The key alias format of the KMS key isn't supported.
- For directory buckets, if you use `PutBucketEncryption` to set your [default bucket encryption](#) to SSE-KMS, Amazon S3 validates the KMS key ID provided in `PutBucketEncryption` requests.

Important

If you're specifying a customer managed KMS key, we recommend using a fully qualified KMS key ARN. If you use a KMS key alias instead, then AWS KMS resolves the key within the requester's account. This behavior can result in data that's encrypted with a KMS key that belongs to the requester, and not the bucket owner.

Also, this action requires AWS Signature Version 4. For more information, see [Authenticating Requests \(AWS Signature Version 4\)](#).

Permissions

- **General purpose bucket permissions** - The `s3:PutEncryptionConfiguration` permission is required in a policy. The bucket owner has this permission by default. The bucket owner can grant this permission to others. For more information about permissions, see [Permissions Related to Bucket Operations](#) and [Managing Access Permissions to Your Amazon S3 Resources](#) in the *Amazon S3 User Guide*.
- **Directory bucket permissions** - To grant access to this API operation, you must have the `s3express:PutEncryptionConfiguration` permission in an IAM identity-based policy instead of a bucket policy. Cross-account access to this API operation isn't supported. This operation can only be performed by the AWS account that owns the resource. For more information about directory bucket policies and permissions, see [AWS Identity and Access Management \(IAM\) for S3 Express One Zone](#) in the *Amazon S3 User Guide*.

To set a directory bucket default encryption with SSE-KMS, you must also have the `kms:GenerateDataKey` and the `kms:Decrypt` permissions in IAM identity-based policies and AWS KMS key policies for the target AWS KMS key.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `s3express-control.region-code.amazonaws.com`.

The following operations are related to `PutBucketEncryption`:

- [GetBucketEncryption](#)
- [DeleteBucketEncryption](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /?encryption HTTP/1.1
```

```

Host: Bucket.s3.amazonaws.com
Content-MD5: ContentMD5
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<ServerSideEncryptionConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Rule>
    <ApplyServerSideEncryptionByDefault>
      <KMSMasterKeyID>string</KMSMasterKeyID>
      <SSEAlgorithm>string</SSEAlgorithm>
    </ApplyServerSideEncryptionByDefault>
    <BlockedEncryptionTypes>
      <EncryptionType>string</EncryptionType>
      ...
    </BlockedEncryptionTypes>
    <BucketKeyEnabled>boolean</BucketKeyEnabled>
  </Rule>
  ...
</ServerSideEncryptionConfiguration>

```

URI Request Parameters

The request uses the following URI parameters.

[Bucket](#)

Specifies default encryption for a bucket using server-side encryption with different key options.

Directory buckets - When you use this operation with a directory bucket, you must use path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name`. Virtual-hosted-style requests aren't supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must also follow the format `bucket-base-name--zone-id--x-s3` (for example, `DOC-EXAMPLE-BUCKET--usw2-az1--x-s3`). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*

Required: Yes

[Content-MD5](#)

The Base64 encoded 128-bit MD5 digest of the server-side encryption configuration.

For requests made using the AWS Command Line Interface (CLI) or AWS SDKs, this field is calculated automatically.

Note

This functionality is not supported for directory buckets.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Note

For directory buckets, this header is not supported in this API operation. If you specify this header, the request fails with the HTTP status code 501 Not Implemented.

x-amz-sdk-checksum-algorithm

Indicates the algorithm used to create the checksum for the request when you use the SDK. This header will not provide any additional functionality if you don't use the SDK. When you send this header, there must be a corresponding `x-amz-checksum` or `x-amz-trailer` header sent. Otherwise, Amazon S3 fails the request with the HTTP status code 400 Bad Request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

If you provide an individual checksum, Amazon S3 ignores any provided `ChecksumAlgorithm` parameter.

Note

For directory buckets, when you use AWS SDKs, CRC32 is the default checksum algorithm that's used for performance.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 | XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in XML format.

[ServerSideEncryptionConfiguration](#)

Root level tag for the ServerSideEncryptionConfiguration parameters.

Required: Yes

[Rule](#)

Container for information about a particular server-side encryption configuration rule.

Type: Array of [ServerSideEncryptionRule](#) data types

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

In the request, you specify the encryption configuration in the request body. The encryption configuration is specified as XML, as shown in the following examples that show setting encryption using SSE-S3, SSE-KMS, or DSSE-KMS.

Request Body for Setting SSE-S3 for general purpose buckets

This example illustrates one usage of PutBucketEncryption.

```
<ServerSideEncryptionConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Rule>
    <ApplyServerSideEncryptionByDefault>
      <SSEAlgorithm>AES256</SSEAlgorithm>
    </ApplyServerSideEncryptionByDefault>
  </Rule>
</ServerSideEncryptionConfiguration>
```

```

    </Rule>
  </ServerSideEncryptionConfiguration>

```

Request Body for Setting SSE-KMS for general purpose buckets

This example illustrates one usage of PutBucketEncryption.

```

<ServerSideEncryptionConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Rule>
    <ApplyServerSideEncryptionByDefault>
      <SSEAlgorithm>aws:kms:dsse</SSEAlgorithm>
      <KMSKeyID>arn:aws:kms:us-east-1:1234/5678example</KMSKeyID>
    </ApplyServerSideEncryptionByDefault>
  </Rule>
</ServerSideEncryptionConfiguration>

```

Set the Default Encryption Configuration for an S3 general purpose bucket

The following is an example of a PUT /? encryption request that specifies to use SSE-KMS encryption.

```

PUT /?encryption HTTP/1.1
Host: examplebucket.<Region>s3.amazonaws.com
Date: Wed, 06 Sep 2017 12:00:00 GMT
Authorization: authorization
Content-Length: length

<ServerSideEncryptionConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Rule>
    <ApplyServerSideEncryptionByDefault>
      <SSEAlgorithm>aws:kms</SSEAlgorithm>
      <KMSKeyID>arn:aws:kms:us-east-1:1234/5678example</KMSKeyID>
    </ApplyServerSideEncryptionByDefault>
  </Rule>
</ServerSideEncryptionConfiguration>

```

Block SSE-C encryption for a bucket

The following is an example of a request to prevent writing new objects to the bucket that uses server-side encryption with customer-provided keys (SSE-C). For more information, see [Blocking or unblocking SSE-C for a general purpose bucket](#).

Note

To allow the use of SSE-C encryption in your write requests to a bucket, pass the value `NONE` instead of `SSE-C`.

```
<ServerSideEncryptionConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Rule>
    <BlockedEncryptionTypes>
      <EncryptionType>SSE-C</EncryptionType>
    </BlockedEncryptionTypes>
  </Rule>
</ServerSideEncryptionConfiguration>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutBucketIntelligentTieringConfiguration

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Puts a S3 Intelligent-Tiering configuration to the specified bucket. You can have up to 1,000 S3 Intelligent-Tiering configurations per bucket.

The S3 Intelligent-Tiering storage class is designed to optimize storage costs by automatically moving data to the most cost-effective storage access tier, without performance impact or operational overhead. S3 Intelligent-Tiering delivers automatic cost savings in three low latency and high throughput access tiers. To get the lowest storage cost on data that can be accessed in minutes to hours, you can choose to activate additional archiving capabilities.

The S3 Intelligent-Tiering storage class is the ideal storage class for data with unknown, changing, or unpredictable access patterns, independent of object size or retention period. If the size of an object is less than 128 KB, it is not monitored and not eligible for auto-tiering. Smaller objects can be stored, but they are always charged at the Frequent Access tier rates in the S3 Intelligent-Tiering storage class.

For more information, see [Storage class for automatically optimizing frequently and infrequently accessed objects](#).

Operations related to PutBucketIntelligentTieringConfiguration include:

- [DeleteBucketIntelligentTieringConfiguration](#)
- [GetBucketIntelligentTieringConfiguration](#)
- [ListBucketIntelligentTieringConfigurations](#)

Note

You only need S3 Intelligent-Tiering enabled on a bucket if you want to automatically move objects stored in the S3 Intelligent-Tiering storage class to the Archive Access or Deep Archive Access tier.

PutBucketIntelligentTieringConfiguration has the following special errors:

HTTP 400 Bad Request Error

Code: InvalidArgument

Cause: Invalid Argument

HTTP 400 Bad Request Error

Code: TooManyConfigurations

Cause: You are attempting to create a new configuration but have already reached the 1,000-configuration limit.

HTTP 403 Forbidden Error

Cause: You are not the owner of the specified bucket, or you do not have the `s3:PutIntelligentTieringConfiguration` bucket permission to set the configuration on the bucket.

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /?intelligent-tiering&id=Id HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<IntelligentTieringConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Id>string</Id>
  <Filter>
    <And>
      <Prefix>string</Prefix>
      <Tag>
        <Key>string</Key>
```

```

        <Value>string</Value>
    </Tag>
    ...
</And>
<Prefix>string</Prefix>
<Tag>
    <Key>string</Key>
    <Value>string</Value>
</Tag>
</Filter>
<Status>string</Status>
<Tiering>
    <AccessTier>string</AccessTier>
    <Days>integer</Days>
</Tiering>
    ...
</IntelligentTieringConfiguration>

```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the Amazon S3 bucket whose configuration you want to modify or retrieve.

Required: Yes

id

The ID used to identify the S3 Intelligent-Tiering configuration.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request accepts the following data in XML format.

[IntelligentTieringConfiguration](#)

Root level tag for the IntelligentTieringConfiguration parameters.

Required: Yes

[Filter](#)

Specifies a bucket filter. The configuration only includes objects that meet the filter's criteria.

Type: [IntelligentTieringFilter](#) data type

Required: No

[Id](#)

The ID used to identify the S3 Intelligent-Tiering configuration.

Type: String

Required: Yes

[Status](#)

Specifies the status of the configuration.

Type: String

Valid Values: Enabled | Disabled

Required: Yes

[Tiering](#)

Specifies the S3 Intelligent-Tiering storage class tier of the configuration.

Type: Array of [Tiering](#) data types

Required: Yes

Response Syntax

```
HTTP/1.1 200
```


Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutBucketInventoryConfiguration

Service: Amazon S3

This implementation of the PUT action adds an S3 Inventory configuration (identified by the inventory ID) to the bucket. You can have up to 1,000 inventory configurations per bucket.

Amazon S3 inventory generates inventories of the objects in the bucket on a daily or weekly basis, and the results are published to a flat file. The bucket that is inventoried is called the *source* bucket, and the bucket where the inventory flat file is stored is called the *destination* bucket. The *destination* bucket must be in the same AWS Region as the *source* bucket.

When you configure an inventory for a *source* bucket, you specify the *destination* bucket where you want the inventory to be stored, and whether to generate the inventory daily or weekly. You can also configure what object metadata to include and whether to inventory all object versions or only current versions. For more information, see [Amazon S3 Inventory](#) in the Amazon S3 User Guide.

Important

You must create a bucket policy on the *destination* bucket to grant permissions to Amazon S3 to write objects to the bucket in the defined location. For an example policy, see [Granting Permissions for Amazon S3 Inventory and Storage Class Analysis](#).

Note

Directory buckets - For directory buckets, you must make requests for this API operation to the Regional endpoint. These endpoints support path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name` . Virtual-hosted-style requests aren't supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

Permissions

To use this operation, you must have permission to perform the `s3:PutInventoryConfiguration` action. The bucket owner has this permission by default and can grant this permission to others.

The `s3:PutInventoryConfiguration` permission allows a user to create an [S3 Inventory](#) report that includes all object metadata fields available and to specify the destination bucket to store the inventory. A user with read access to objects in the destination bucket can also access all object metadata fields that are available in the inventory report.

- **General purpose bucket permissions** - The `s3:PutInventoryConfiguration` permission is required in a policy. For more information about general purpose buckets permissions, see [Using Bucket Policies and User Policies](#) in the *Amazon S3 User Guide*.
- **Directory bucket permissions** - To grant access to this API operation, you must have the `s3express:PutInventoryConfiguration` permission in an IAM identity-based policy instead of a bucket policy. For more information about directory bucket policies and permissions, see [AWS Identity and Access Management \(IAM\) for S3 Express One Zone](#) in the *Amazon S3 User Guide*.

To restrict access to an inventory report, see [Restricting access to an Amazon S3 Inventory report](#) in the *Amazon S3 User Guide*. For more information about the metadata fields available in S3 Inventory, see [Amazon S3 Inventory lists](#) in the *Amazon S3 User Guide*. For more information about permissions, see [Permissions related to bucket subresource operations](#) and [Identity and access management in Amazon S3](#) in the *Amazon S3 User Guide*.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `s3express-control.region-code.amazonaws.com`.

`PutBucketInventoryConfiguration` has the following special errors:

HTTP 400 Bad Request Error

Code: `InvalidArgument`

Cause: Invalid Argument

HTTP 400 Bad Request Error

Code: `TooManyConfigurations`

Cause: You are attempting to create a new configuration but have already reached the 1,000-configuration limit.

HTTP 403 Forbidden Error

Cause: You are not the owner of the specified bucket, or you do not have the `s3:PutInventoryConfiguration` bucket permission to set the configuration on the bucket.

The following operations are related to `PutBucketInventoryConfiguration`:

- [GetBucketInventoryConfiguration](#)
- [DeleteBucketInventoryConfiguration](#)
- [ListBucketInventoryConfigurations](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /?inventory&id=Id HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<InventoryConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Destination>
    <S3BucketDestination>
      <AccountId>string</AccountId>
      <Bucket>string</Bucket>
      <Encryption>
        <SSE-KMS>
          <KeyId>string</KeyId>
        </SSE-KMS>
        <SSE-S3>
          </SSE-S3>
        </Encryption>
      <Format>string</Format>
      <Prefix>string</Prefix>
    </S3BucketDestination>
  </Destination>
</InventoryConfiguration>
```

```

</Destination>
<IsEnabled>boolean</IsEnabled>
<Filter>
  <Prefix>string</Prefix>
</Filter>
<Id>string</Id>
<IncludedObjectVersions>string</IncludedObjectVersions>
<OptionalFields>
  <Field>string</Field>
</OptionalFields>
<Schedule>
  <Frequency>string</Frequency>
</Schedule>
</InventoryConfiguration>

```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket where the inventory configuration will be stored.

Directory buckets - When you use this operation with a directory bucket, you must use path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name` . Virtual-hosted-style requests aren't supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must also follow the format `bucket-base-name--zone-id--x-s3` (for example, `DOC-EXAMPLE-BUCKET--usw2-az1--x-s3`). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*

Required: Yes

id

The ID used to identify the inventory configuration.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Note

For directory buckets, this header is not supported in this API operation. If you specify this header, the request fails with the HTTP status code 501 Not Implemented.

Request Body

The request accepts the following data in XML format.

InventoryConfiguration

Root level tag for the InventoryConfiguration parameters.

Required: Yes

Destination

Contains information about where to publish the inventory results.

Type: [InventoryDestination](#) data type

Required: Yes

Filter

Specifies an inventory filter. The inventory only includes objects that meet the filter's criteria.

Type: [InventoryFilter](#) data type

Required: No

Id

The ID used to identify the inventory configuration.

Type: String

Required: Yes

IncludedObjectVersions

Object versions to include in the inventory list. If set to `All`, the list includes all the object versions, which adds the version-related fields `VersionId`, `IsLatest`, and `DeleteMarker` to the list. If set to `Current`, the list does not contain these version-related fields.

Type: String

Valid Values: All | Current

Required: Yes

IsEnabled

Specifies whether the inventory is enabled or disabled. If set to `True`, an inventory list is generated. If set to `False`, no inventory list is generated.

Type: Boolean

Required: Yes

OptionalFields

Contains the optional fields that are included in the inventory results.

Note

The following optional fields are supported for directory buckets `Size` | `LastModifiedDate` | `StorageClass` | `ETag` | `IsMultipartUploaded` | `EncryptionStatus` | `BucketKeyStatus` | `ChecksumAlgorithm` | `LifecycleExpirationDate`. Throws `MalformedXML` error if unsupported optional field is provided.

Type: Array of strings

Valid Values: `Size` | `LastModifiedDate` | `StorageClass` | `ETag` | `IsMultipartUploaded` | `ReplicationStatus` | `EncryptionStatus` | `ObjectLockRetainUntilDate` | `ObjectLockMode` | `ObjectLockLegalHoldStatus` | `ObjectLockEventHoldStatus` | `ObjectLockEventHoldDuration` | `IntelligentTieringAccessTier` | `BucketKeyStatus` | `ChecksumAlgorithm` | `ObjectAccessControlList` | `ObjectOwner` | `LifecycleExpirationDate`

Required: No

Schedule

Specifies the schedule for generating inventory results.

Type: [InventorySchedule](#) data type

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Example: Create an inventory configuration

The following PUT request and response for the bucket `examplebucket` creates a new or replaces an existing inventory configuration with the ID `report1`. The configuration is defined in the request body.

```
PUT /?inventory&id=report1 HTTP/1.1
Host: examplebucket.s3.<Region>.amazonaws.com
Date: Mon, 31 Oct 2016 12:00:00 GMT
Authorization: authorization string
Content-Length: length

<?xml version="1.0" encoding="UTF-8"?>
<InventoryConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Id>report1</Id>
  <IsEnabled>true</IsEnabled>
  <Filter>
    <Prefix>filterPrefix</Prefix>
  </Filter>
  <Destination>
    <S3BucketDestination>
      <Format>CSV</Format>
      <AccountId>123456789012</AccountId>
      <Bucket>arn:aws:s3:::destination-bucket</Bucket>
      <Prefix>prefix1</Prefix>
      <Encryption>
        <SSE-KMS>
```



```

        <KeyId>arn:aws:kms:us-
west-2:111122223333:key/1234abcd-12ab-34cd-56ef-1234567890ab</KeyId>
    </SSE-KMS>
</Encryption>
</S3BucketDestination>
</Destination>
<Schedule>
    <Frequency>Daily</Frequency>
</Schedule>
<IncludedObjectVersions>All</IncludedObjectVersions>
<OptionalFields>
    <Field>Size</Field>
    <Field>LastModifiedDate</Field>
    <Field>ETag</Field>
    <Field>StorageClass</Field>
    <Field>IsMultipartUploaded</Field>
    <Field>ReplicationStatus</Field>
    <Field>EncryptionStatus</Field>
    <Field>ObjectLockRetainUntilDate</Field>
    <Field>ObjectLockMode</Field>
    <Field>ObjectLockLegalHoldStatus</Field>
</OptionalFields>
</InventoryConfiguration>

```

```

HTTP/1.1 200 OK
x-amz-id-2: YgIPIfBiKa2bj0KMg95r/0zo3emzU4dzsD4rcKCHQUAdQkf3ShJT00pXUueF6QKo
x-amz-request-id: 236A8905248E5A01
Date: Mon, 31 Oct 2016 12:00:00 GMT
Content-Length: 0
Server: AmazonS3

```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)

- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutBucketLifecycle

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Important

For an updated version of this API, see [PutBucketLifecycleConfiguration](#). This version has been deprecated. Existing lifecycle configurations will work. For new lifecycle configurations, use the updated API.

Note

This operation is not supported for directory buckets.

Creates a new lifecycle configuration for the bucket or replaces an existing lifecycle configuration. For information about lifecycle configuration, see [Object Lifecycle Management](#) in the *Amazon S3 User Guide*.

By default, all Amazon S3 resources, including buckets, objects, and related subresources (for example, lifecycle configuration and website configuration) are private. Only the resource owner, the AWS account that created the resource, can access it. The resource owner can optionally grant access permissions to others by writing an access policy. For this operation, users must get the `s3:PutLifecycleConfiguration` permission.

You can also explicitly deny permissions. Explicit denial also supersedes any other permissions. If you want to prevent users or accounts from removing or deleting objects from your bucket, you must deny them permissions for the following actions:

- `s3:DeleteObject`
- `s3:DeleteObjectVersion`
- `s3:PutLifecycleConfiguration`

For more information about permissions, see [Managing Access Permissions to your Amazon S3 Resources](#) in the *Amazon S3 User Guide*.

For more examples of transitioning objects to storage classes such as STANDARD_IA or ONEZONE_IA, see [Examples of Lifecycle Configuration](#).

The following operations are related to PutBucketLifecycle:

- [GetBucketLifecycle](#)(Deprecated)
- [GetBucketLifecycleConfiguration](#)
- [RestoreObject](#)
- By default, a resource owner—in this case, a bucket owner, which is the AWS account that created the bucket—can perform any of the operations. A resource owner can also grant others permission to perform the operation. For more information, see the following topics in the Amazon S3 User Guide:
 - [Specifying Permissions in a Policy](#)
 - [Managing Access Permissions to your Amazon S3 Resources](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /?lifecycle HTTP/1.1
Host: Bucket.s3.amazonaws.com
Content-MD5: ContentMD5
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<LifecycleConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Rule>
    <AbortIncompleteMultipartUpload>
      <DaysAfterInitiation>integer</DaysAfterInitiation>
    </AbortIncompleteMultipartUpload>
```

```

<Expiration>
  <Date>timestamp</Date>
  <Days>integer</Days>
  <ExpiredObjectDeleteMarker>boolean</ExpiredObjectDeleteMarker>
</Expiration>
<ID>string</ID>
<NoncurrentVersionExpiration>
  <NewerNoncurrentVersions>integer</NewerNoncurrentVersions>
  <NoncurrentDays>integer</NoncurrentDays>
</NoncurrentVersionExpiration>
<NoncurrentVersionTransition>
  <NewerNoncurrentVersions>integer</NewerNoncurrentVersions>
  <NoncurrentDays>integer</NoncurrentDays>
  <StorageClass>string</StorageClass>
</NoncurrentVersionTransition>
<Prefix>string</Prefix>
<Status>string</Status>
<Transition>
  <Date>timestamp</Date>
  <Days>integer</Days>
  <StorageClass>string</StorageClass>
</Transition>
</Rule>
...
</LifecycleConfiguration>

```

URI Request Parameters

The request uses the following URI parameters.

Bucket

Required: Yes

Content-MD5

For requests made using the AWS Command Line Interface (CLI) or AWS SDKs, this field is calculated automatically.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-sdk-checksum-algorithm

Indicates the algorithm used to create the checksum for the request when you use the SDK. This header will not provide any additional functionality if you don't use the SDK. When you send this header, there must be a corresponding x-amz-checksum or x-amz-trailer header sent. Otherwise, Amazon S3 fails the request with the HTTP status code 400 Bad Request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

If you provide an individual checksum, Amazon S3 ignores any provided ChecksumAlgorithm parameter.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 | XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in XML format.

LifecycleConfiguration

Root level tag for the LifecycleConfiguration parameters.

Required: Yes

Rule

Specifies lifecycle configuration rules for an Amazon S3 bucket.

Type: Array of [Rule](#) data types

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Sample Request: Body of a basic lifecycle configuration

In the request, you specify the lifecycle configuration in the request body. The lifecycle configuration is specified as XML. The following is an example of a basic lifecycle configuration. It specifies one rule. The `Prefix` in the rule identifies objects to which the rule applies. The rule also specifies two actions (`Transition` and `Expiration`). Each action specifies a time line when Amazon S3 should perform the action. The `Status` indicates whether the rule is enabled or disabled.

```
<LifecycleConfiguration>
  <Rule>
    <ID>sample-rule</ID>
    <Prefix>key-prefix</Prefix>
    <Status>rule-status</Status>
    <Transition>
      <Date>value</Date>
      <StorageClass>storage class</StorageClass>
    </Transition>
    <Expiration>
      <Days>value</Days>
    </Expiration>
  </Rule>
</LifecycleConfiguration>
```

Sample Request: Body of a lifecycle configuration specifying noncurrent versions

If the state of your bucket is versioning-enabled or versioning-suspended, you can have many versions of the same object: one current version and zero or more noncurrent versions. The following lifecycle configuration specifies the actions (`NoncurrentVersionTransition`, `NoncurrentVersionExpiration`) that are specific to noncurrent object versions.

```
<LifecycleConfiguration>
  <Rule>
    <ID>sample-rule</ID>
    <Prefix>key-prefix</Prefix>
    <Status>rule-status</Status>
    <NoncurrentVersionTransition>
      <NoncurrentDays>value</NoncurrentDays>
      <StorageClass>storage class</StorageClass>
    </NoncurrentVersionTransition>
    <NoncurrentVersionExpiration>
      <NoncurrentDays>value</NoncurrentDays>
    </NoncurrentVersionExpiration>
  </Rule>
</LifecycleConfiguration>
```

Sample Request: Body of a lifecycle configuration that specifies a rule with `AbortIncompleteMultipartUpload`

You can use the multipart upload to upload large objects in parts. For more information about multipart uploads, see [Multipart Upload Overview](#) in the *Amazon S3 User Guide*. With lifecycle configuration, you can tell Amazon S3 to abort incomplete multipart uploads, which are identified by the key name prefix specified in the rule, if they don't complete within a specified number of days. When Amazon S3 aborts a multipart upload, it deletes all parts associated with the upload. This ensures that you don't have incomplete multipart uploads that have left parts stored in Amazon S3, so you don't have to pay storage costs for them. The following is an example lifecycle configuration that specifies a rule with the `AbortIncompleteMultipartUpload` action. This action tells Amazon S3 to abort incomplete multipart uploads seven days after initiation.

```
<LifecycleConfiguration>
  <Rule>
    <ID>sample-rule</ID>
    <Prefix>SomeKeyPrefix</Prefix>
    <Status>rule-status</Status>
    <AbortIncompleteMultipartUpload>
      <DaysAfterInitiation>7</DaysAfterInitiation>
    </AbortIncompleteMultipartUpload>
  </Rule>
</LifecycleConfiguration>
```


Add lifecycle configuration to a bucket that is not versioning-enabled

The following is a sample PUT `/?lifecycle` request that adds the lifecycle configuration to the `examplebucket` bucket. The lifecycle configuration specifies two rules, each with one action:

- The **Transition** action tells Amazon S3 to transition objects with the "documents/" prefix to the GLACIER storage class 30 days after creation.
- The **Expiration** action tells Amazon S3 to delete objects with the "logs/" prefix 365 days after creation.

The sample response follows the sample request.

```
PUT /?lifecycle HTTP/1.1
Host: examplebucket.s3.<Region>.amazonaws.com
x-amz-date: Wed, 14 May 2014 02:11:21 GMT
Content-MD5: q6yJDlIkBaGGfb3QLY69A==
Authorization: authorization string
Content-Length: 415
<LifecycleConfiguration>
  <Rule>
    <ID>id1</ID>
    <Prefix>documents</Prefix>
    <Status>Enabled</Status>
    <Transition>
      <Days>30</Days>
      <StorageClass>GLACIER</StorageClass>
    </Transition>
  </Rule>
  <Rule>
    <ID>id2</ID>
    <Prefix>logs</Prefix>
    <Status>Enabled</Status>
    <Expiration>
      <Days>365</Days>
    </Expiration>
  </Rule>
</LifecycleConfiguration>
```

```
HTTP/1.1 200 OK
x-amz-id-2: r+qR7+nhXtJDDIJ0JJYcd+1j5nM/rUFiiiZ/fNbD0sd3JUE8NWMLNHXmvPfwMpdC
```

```
x-amz-request-id: 9E26D08072A8EF9E
Date: Wed, 14 May 2014 02:11:22 GMT
Content-Length: 0
Server: AmazonS3
```

Add lifecycle configuration to a bucket that is versioning-enabled

The following is a sample PUT `/?lifecycle` request that adds the lifecycle configuration to the `examplebucket` bucket. The lifecycle configuration specifies two rules, each with one action. You specify these actions when your bucket is versioning-enabled or versioning is suspended:

- The `NoncurrentVersionExpiration` action tells Amazon S3 to expire noncurrent versions of objects with the "logs/" prefix 100 days after the objects become noncurrent.
- The `NoncurrentVersionTransition` action tells Amazon S3 to transition noncurrent versions of objects with the "documents/" prefix to the GLACIER storage class 30 days after they become noncurrent.

The sample response follows the sample request.

```
PUT /?lifecycle HTTP/1.1
Host: examplebucket.s3.<Region>.amazonaws.com
x-amz-date: Wed, 14 May 2014 02:21:48 GMT
Content-MD5: 96rxH9mDqVnKkaZDddgnw==
Authorization: authorization string
Content-Length: 598
<LifecycleConfiguration>
  <Rule>
    <ID>id1</ID>
    <Prefix>logs/</Prefix>
    <Status>Enabled</Status>
    <NoncurrentVersionExpiration>
      <NoncurrentDays>1</NoncurrentDays>
    </NoncurrentVersionExpiration>
  </Rule>
  <Rule>
    <ID>TransitionSoonAfterBecomingNonCurrent</ID>
    <Prefix>documents/</Prefix>
    <Status>Enabled</Status>
    <NoncurrentVersionTransition>
      <NoncurrentDays>0</NoncurrentDays>
```

```
<StorageClass>GLACIER</StorageClass>
</NoncurrentVersionTransition>
</Rule>
</LifecycleConfiguration>
```

```
HTTP/1.1 200 OK
x-amz-id-2: aXQ+KbIrmMmo0//3bMdDTw/CnjArwje+J49Hf+j44yRb/VmbIkgI05A+PT98Cp/6k07hf
+LD2mY=
x-amz-request-id: 02D7EC4C10381EB1
Date: Wed, 14 May 2014 02:21:50 GMT
Content-Length: 0
Server: AmazonS3
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutBucketLifecycleConfiguration

Service: Amazon S3

Creates a new lifecycle configuration for the bucket or replaces an existing lifecycle configuration. Keep in mind that this will overwrite an existing lifecycle configuration, so if you want to retain any configuration details, they must be included in the new lifecycle configuration. For information about lifecycle configuration, see [Managing your storage lifecycle](#).

Note

Bucket lifecycle configuration now supports specifying a lifecycle rule using an object key name prefix, one or more object tags, object size, or any combination of these. Accordingly, this section describes the latest API. The previous version of the API supported filtering based only on an object key name prefix, which is supported for backward compatibility. For the related API description, see [PutBucketLifecycle](#).

Rules

You specify the lifecycle configuration in your request body. The lifecycle configuration is specified as XML consisting of one or more rules. An Amazon S3 Lifecycle configuration can have up to 1,000 rules. This limit is not adjustable.

Bucket lifecycle configuration supports specifying a lifecycle rule using an object key name prefix, one or more object tags, object size, or any combination of these. Accordingly, this section describes the latest API. The previous version of the API supported filtering based only on an object key name prefix, which is supported for backward compatibility for general purpose buckets. For the related API description, see [PutBucketLifecycle](#).

Note

Lifecycle configurations for directory buckets only support expiring objects and cancelling multipart uploads. Expiring of versioned objects, transitions and tag filters are not supported.

A lifecycle rule consists of the following:

- A filter identifying a subset of objects to which the rule applies. The filter can be based on a key name prefix, object tags, object size, or any combination of these.

- A status indicating whether the rule is in effect.
- One or more lifecycle transition and expiration actions that you want Amazon S3 to perform on the objects identified by the filter. If the state of your bucket is versioning-enabled or versioning-suspended, you can have many versions of the same object (one current version and zero or more noncurrent versions). Amazon S3 provides predefined actions that you can specify for current and noncurrent object versions.

For more information, see [Object Lifecycle Management](#) and [Lifecycle Configuration Elements](#).

Permissions

- **General purpose bucket permissions** - By default, all Amazon S3 resources are private, including buckets, objects, and related subresources (for example, lifecycle configuration and website configuration). Only the resource owner (that is, the AWS account that created it) can access the resource. The resource owner can optionally grant access permissions to others by writing an access policy. For this operation, a user must have the `s3:PutLifecycleConfiguration` permission.

You can also explicitly deny permissions. An explicit deny also supersedes any other permissions. If you want to block users or accounts from removing or deleting objects from your bucket, you must deny them permissions for the following actions:

- `s3:DeleteObject`
- `s3:DeleteObjectVersion`
- `s3:PutLifecycleConfiguration`

For more information about permissions, see [Managing Access Permissions to Your Amazon S3 Resources](#).

- **Directory bucket permissions** - You must have the `s3express:PutLifecycleConfiguration` permission in an IAM identity-based policy to use this operation. Cross-account access to this API operation isn't supported. The resource owner can optionally grant access permissions to others by creating a role or user for them as long as they are within the same account as the owner and resource.

For more information about directory bucket policies and permissions, see [Authorizing Regional endpoint APIs with IAM](#) in the *Amazon S3 User Guide*.

Note

Directory buckets - For directory buckets, you must make requests for this API operation to the Regional endpoint. These endpoints support path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name`. Virtual-hosted-style requests aren't supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `s3express-control.region.amazonaws.com`.

The following operations are related to `PutBucketLifecycleConfiguration`:

- [GetBucketLifecycleConfiguration](#)
- [DeleteBucketLifecycle](#)

⚠ Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /?lifecycle HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-expected-bucket-owner: ExpectedBucketOwner
x-amz-transition-default-minimum-object-size: TransitionDefaultMinimumObjectSize
<?xml version="1.0" encoding="UTF-8"?>
<LifecycleConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Rule>
    <AbortIncompleteMultipartUpload>
```

```

    <DaysAfterInitiation>integer</DaysAfterInitiation>
  </AbortIncompleteMultipartUpload>
  <Expiration>
    <Date>timestamp</Date>
    <Days>integer</Days>
    <ExpiredObjectDeleteMarker>boolean</ExpiredObjectDeleteMarker>
  </Expiration>
  <Filter>
    <And>
      <ObjectSizeGreaterThan>long</ObjectSizeGreaterThan>
      <ObjectSizeLessThan>long</ObjectSizeLessThan>
      <Prefix>string</Prefix>
      <Tag>
        <Key>string</Key>
        <Value>string</Value>
      </Tag>
      ...
    </And>
    <ObjectSizeGreaterThan>long</ObjectSizeGreaterThan>
    <ObjectSizeLessThan>long</ObjectSizeLessThan>
    <Prefix>string</Prefix>
    <Tag>
      <Key>string</Key>
      <Value>string</Value>
    </Tag>
  </Filter>
  <ID>string</ID>
  <NoncurrentVersionExpiration>
    <NewerNoncurrentVersions>integer</NewerNoncurrentVersions>
    <NoncurrentDays>integer</NoncurrentDays>
  </NoncurrentVersionExpiration>
  <NoncurrentVersionTransition>
    <NewerNoncurrentVersions>integer</NewerNoncurrentVersions>
    <NoncurrentDays>integer</NoncurrentDays>
    <StorageClass>string</StorageClass>
  </NoncurrentVersionTransition>
  ...
  <Prefix>string</Prefix>
  <Status>string</Status>
  <Transition>
    <Date>timestamp</Date>
    <Days>integer</Days>
    <StorageClass>string</StorageClass>
  </Transition>

```

```
...  
</Rule>  
...  
</LifecycleConfiguration>
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket for which to set the configuration.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Note

This parameter applies to general purpose buckets only. It is not supported for directory bucket lifecycle configurations.

x-amz-sdk-checksum-algorithm

Indicates the algorithm used to create the checksum for the request when you use the SDK. This header will not provide any additional functionality if you don't use the SDK. When you send this header, there must be a corresponding `x-amz-checksum` or `x-amz-trailer` header sent. Otherwise, Amazon S3 fails the request with the HTTP status code 400 Bad Request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

If you provide an individual checksum, Amazon S3 ignores any provided `ChecksumAlgorithm` parameter.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 |
XXHASH64 | XXHASH3 | XXHASH128

x-amz-transition-default-minimum-object-size

Indicates which default minimum object size behavior is applied to the lifecycle configuration.

Note

This parameter applies to general purpose buckets only. It is not supported for directory bucket lifecycle configurations.

- `all_storage_classes_128K` - Objects smaller than 128 KB will not transition to any storage class by default.
- `varies_by_storage_class` - Objects smaller than 128 KB will transition to Glacier Flexible Retrieval or Glacier Deep Archive storage classes. By default, all other storage classes will prevent transitions smaller than 128 KB.

To customize the minimum object size for any transition you can add a filter that specifies a custom `ObjectSizeGreaterThan` or `ObjectSizeLessThan` in the body of your transition rule. Custom filters always take precedence over the default transition behavior.

Valid Values: `varies_by_storage_class` | `all_storage_classes_128K`

Request Body

The request accepts the following data in XML format.

LifecycleConfiguration

Root level tag for the LifecycleConfiguration parameters.

Required: Yes

Rule

A lifecycle rule for individual objects in an Amazon S3 bucket.

Type: Array of [LifecycleRule](#) data types

Required: Yes

Response Syntax

```
HTTP/1.1 200
x-amz-transition-default-minimum-object-size: TransitionDefaultMinimumObjectSize
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

[x-amz-transition-default-minimum-object-size](#)

Indicates which default minimum object size behavior is applied to the lifecycle configuration.

Note

This parameter applies to general purpose buckets only. It is not supported for directory bucket lifecycle configurations.

- `all_storage_classes_128K` - Objects smaller than 128 KB will not transition to any storage class by default.
- `varies_by_storage_class` - Objects smaller than 128 KB will transition to Glacier Flexible Retrieval or Glacier Deep Archive storage classes. By default, all other storage classes will prevent transitions smaller than 128 KB.

To customize the minimum object size for any transition you can add a filter that specifies a custom `ObjectSizeGreaterThan` or `ObjectSizeLessThan` in the body of your transition rule. Custom filters always take precedence over the default transition behavior.

Valid Values: `varies_by_storage_class` | `all_storage_classes_128K`

Examples

Example 1: Add lifecycle configuration - bucket not versioning-enabled

The following lifecycle configuration for a general purpose bucket specifies two rules, each with one action.

- The Transition action requests Amazon S3 to transition objects with the "documents/" prefix to the GLACIER storage class 30 days after creation.
- The Expiration action requests Amazon S3 to delete objects with the "logs/" prefix 365 days after creation.

```
<LifecycleConfiguration>
  <Rule>
    <ID>id1</ID>
    <Filter>
      <Prefix>documents/</Prefix>
    </Filter>
    <Status>Enabled</Status>
    <Transition>
      <Days>30</Days>
      <StorageClass>GLACIER</StorageClass>
    </Transition>
  </Rule>
  <Rule>
    <ID>id2</ID>
    <Filter>
      <Prefix>logs/</Prefix>
    </Filter>
    <Status>Enabled</Status>
    <Expiration>
      <Days>365</Days>
    </Expiration>
  </Rule>
</LifecycleConfiguration>
```

Example 2: Adding a lifecycle configuration to a general purpose bucket

The following is a sample PUT `/?lifecycle` request that adds the preceding lifecycle configuration to a general purpose bucket.

```
PUT /?lifecycle HTTP/1.1
Host: examplebucket.s3.<Region>.amazonaws.com
x-amz-date: Wed, 14 May 2014 02:11:21 GMT
```

```
Content-MD5: q6yJDlIkBaGGfb3QLY69A==
Authorization: authorization string
Content-Length: 415

<LifecycleConfiguration>
  <Rule>
    <ID>id1</ID>
    <Filter>
      <Prefix>documents/</Prefix>
    </Filter>
    <Status>Enabled</Status>
    <Transition>
      <Days>30</Days>
      <StorageClass>GLACIER</StorageClass>
    </Transition>
  </Rule>
  <Rule>
    <ID>id2</ID>
    <Filter>
      <Prefix>logs/</Prefix>
    </Filter>
    <Status>Enabled</Status>
    <Expiration>
      <Days>365</Days>
    </Expiration>
  </Rule>
</LifecycleConfiguration>
```

Sample Response

This example illustrates one usage of PutBucketLifecycleConfiguration.

```
HTTP/1.1 200 OK
x-amz-id-2: r+qR7+nhXtJDDIJ0JJYcd+1j5nM/rUFiiiZ/fNbD0sd3JUE8NWMLNHXmvPfwMpdC
x-amz-request-id: 9E26D08072A8EF9E
Date: Wed, 14 May 2014 02:11:22 GMT
Content-Length: 0
Server: AmazonS3
```

Example 3: Add a lifecycle configuration - bucket is versioning-enabled

The following lifecycle configuration for a general purpose bucket specifies two rules, each with one action for Amazon S3 to perform. You specify these actions when your bucket is versioning-enabled or versioning is suspended:

- The `NoncurrentVersionExpiration` action requests Amazon S3 to expire noncurrent versions of objects with the "logs/" prefix 100 days after the objects become noncurrent.
- The `NoncurrentVersionTransition` action requests Amazon S3 to transition noncurrent versions of objects with the "documents/" prefix to the GLACIER storage class 30 days after they become noncurrent.

```
<LifecycleConfiguration>
  <Rule>
    <ID>DeleteAfterBecomingNonCurrent</ID>
    <Filter>
      <Prefix>logs/</Prefix>
    </Filter>
    <Status>Enabled</Status>
    <NoncurrentVersionExpiration>
      <NoncurrentDays>100</NoncurrentDays>
    </NoncurrentVersionExpiration>
  </Rule>
  <Rule>
    <ID>TransitionAfterBecomingNonCurrent</ID>
    <Filter>
      <Prefix>documents/</Prefix>
    </Filter>
    <Status>Enabled</Status>
    <NoncurrentVersionTransition>
      <NoncurrentDays>30</NoncurrentDays>
      <StorageClass>GLACIER</StorageClass>
    </NoncurrentVersionTransition>
  </Rule>
</LifecycleConfiguration>
```

Example 4: Add a lifecycle configuration to a general purpose bucket

The following is a sample PUT `/?lifecycle` request that adds the preceding lifecycle configuration to a `examplebucket` bucket.

```
PUT /?lifecycle HTTP/1.1
Host: examplebucket.s3.<Region>.amazonaws.com
x-amz-date: Wed, 14 May 2014 02:21:48 GMT
Content-MD5: 96rxH9mDqVnKkaZDddgnw==
Authorization: authorization string
Content-Length: 598

<LifecycleConfiguration>
  <Rule>
    <ID>DeleteAfterBecomingNonCurrent</ID>
    <Filter>
      <Prefix>logs/</Prefix>
    </Filter>
    <Status>Enabled</Status>
    <NoncurrentVersionExpiration>
      <NoncurrentDays>1</NoncurrentDays>
    </NoncurrentVersionExpiration>
  </Rule>
  <Rule>
    <ID>TransitionSoonAfterBecomingNonCurrent</ID>
    <Filter>
      <Prefix>documents/</Prefix>
    </Filter>
    <Status>Enabled</Status>
    <NoncurrentVersionTransition>
      <NoncurrentDays>0</NoncurrentDays>
      <StorageClass>GLACIER</StorageClass>
    </NoncurrentVersionTransition>
  </Rule>
</LifecycleConfiguration>
```

Sample Response

This example illustrates one usage of `PutBucketLifecycleConfiguration`.

```

HTTP/1.1 200 OK
x-amz-id-2: aXQ+KbIrmMmo0//3bMdDTw/CnjArwje+J49Hf+j44yRb/VmbIkgI05A+PT98Cp/6k07hf
+LD2mY=
x-amz-request-id: 02D7EC4C10381EB1
Date: Wed, 14 May 2014 02:21:50 GMT
Content-Length: 0
Server: AmazonS3

```

Example 5: Add a lifecycle configuration to a directory bucket

The following lifecycle configuration specifies two rules, each with one action for Amazon S3 to perform:

- The `Expiration` action requests Amazon S3 to expire objects with object size between 500B to 64000B and the "myprefix/" prefix 7 days after their creation.
- The `AbortIncompleteMultipartUpload` action requests Amazon S3 to abort incomplete multipart uploads 3 days after their initiation.

```

PUT /?lifecycle HTTP/1.1
Host: s3express-control.us-west-2.amazonaws.com
x-amz-sdk-checksum-algorithm: CRC32
x-amz-checksum-crc32: UCqxTw==

<LifecycleConfiguration>
  <Rule>
    <ID>Lifecycle expiration rule</ID>
    <Filter>
      <And>
        <Prefix>myprefix/</Prefix>
        <ObjectSizeGreaterThan>500</ObjectSizeGreaterThan>
        <ObjectSizeLessThan>64000</ObjectSizeLessThan>
      </And>
    <Filter>
      <Status>Enabled</Status>
    <Expiration>
      <Days>7</Days>
    <Expiration>
  </Rule>
</Rule>

```

```
<ID>MPU Rule</ID>
<Filter>
  <Prefix>another_prefix</Prefix>
</Filter>
<Status>Enabled</Status>
<AbortIncompleteMultipartUpload>
  <DaysAfterInitiation>3</DaysAfterInitiation>
</AbortIncompleteMultipartUpload>
</Rule>
</LifecycleConfiguration>
```

Sample Response

This example illustrates one usage of `PutBucketLifecycleConfiguration` in directory buckets

```
HTTP/1.1 200 OK
Server: AmazonS3
x-amz-request-id: 02D7EC4C10381EB1
x-amz-id-2: kc3B1Ns0xqzIdJNd
Content-Type: application/xml
Content-Length: 0
Date: Tue, 12 Nov 2024 18:59:45 GMT
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)

- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutBucketLogging

Service: Amazon S3

Important

End of support notice: As of October 1, 2025, Amazon S3 has discontinued support for Email Grantee Access Control Lists (ACLs). If you attempt to use an Email Grantee ACL in a request after October 1, 2025, the request will receive an HTTP 405 (Method Not Allowed) error.

This change affects the following AWS Regions: US East (N. Virginia), US West (N. California), US West (Oregon), Asia Pacific (Singapore), Asia Pacific (Sydney), Asia Pacific (Tokyo), Europe (Ireland), and South America (São Paulo).

Note

This operation is not supported for directory buckets.

Set the logging parameters for a bucket and to specify permissions for who can view and modify the logging parameters. All logs are saved to buckets in the same AWS Region as the source bucket. To set the logging status of a bucket, you must be the bucket owner.

The bucket owner is automatically granted FULL_CONTROL to all logs. You use the Grantee request element to grant access to other people. The Permissions request element specifies the kind of access the grantee has to the logs.

Important

If the target bucket for log delivery uses the bucket owner enforced setting for S3 Object Ownership, you can't use the Grantee request element to grant access to others. Permissions can only be granted using policies. For more information, see [Permissions for server access log delivery](#) in the *Amazon S3 User Guide*.

Grantee Values

You can specify the person (grantee) to whom you're assigning access rights (by using request elements) in the following ways. For examples of how to specify these grantee values in JSON format, see the AWS CLI example in [Enabling Amazon S3 server access logging](#) in the *Amazon S3 User Guide*.

- By the person's ID:

```
<Grantee xmlns:xsi="http://www.w3.org/2001/XMLSchema-
instance" xsi:type="CanonicalUser"><ID><>ID<></
ID><DisplayName><>GranteesEmail<></DisplayName> </Grantee>
```

DisplayName is optional and ignored in the request.

- By Email address:

```
<Grantee xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:type="AmazonCustomerByEmail"><EmailAddress><>Grantees@email.com<></
EmailAddress></Grantee>
```

The grantee is resolved to the CanonicalUser and, in a response to a GETObjectAcl request, appears as the CanonicalUser.

- By URI:

```
<Grantee xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:type="Group"><URI><>http://acs.amazonaws.com/groups/global/
AuthenticatedUsers<></URI></Grantee>
```

To enable logging, you use LoggingEnabled and its children request elements. To disable logging, you use an empty BucketLoggingStatus request element:

```
<BucketLoggingStatus xmlns="http://doc.s3.amazonaws.com/2006-03-01" />
```

For more information about server access logging, see [Server Access Logging](#) in the *Amazon S3 User Guide*.

For more information about creating a bucket, see [CreateBucket](#). For more information about returning the logging status of a bucket, see [GetBucketLogging](#).

The following operations are related to PutBucketLogging:

- [PutObject](#)
- [DeleteBucket](#)
- [CreateBucket](#)
- [GetBucketLogging](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /?logging HTTP/1.1
Host: Bucket.s3.amazonaws.com
Content-MD5: ContentMD5
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<BucketLoggingStatus xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <LoggingEnabled>
    <TargetBucket>string</TargetBucket>
    <TargetGrants>
      <Grant>
        <Grantee>
          <DisplayName>string</DisplayName>
          <EmailAddress>string</EmailAddress>
          <ID>string</ID>
          <xsi:type>string</xsi:type>
          <URI>string</URI>
        </Grantee>
        <Permission>string</Permission>
      </Grant>
    </TargetGrants>
    <TargetObjectKeyFormat>
      <PartitionedPrefix>
        <PartitionDateSource>string</PartitionDateSource>
      </PartitionedPrefix>
      <SimplePrefix>
```

```
</SimplePrefix>
</TargetObjectKeyFormat>
<TargetPrefix>string</TargetPrefix>
</LoggingEnabled>
</BucketLoggingStatus>
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket for which to set the logging parameters.

Required: Yes

Content-MD5

The MD5 hash of the PutBucketLogging request body.

For requests made using the AWS Command Line Interface (CLI) or AWS SDKs, this field is calculated automatically.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-sdk-checksum-algorithm

Indicates the algorithm used to create the checksum for the request when you use the SDK. This header will not provide any additional functionality if you don't use the SDK. When you send this header, there must be a corresponding x-amz-checksum or x-amz-trailer header sent. Otherwise, Amazon S3 fails the request with the HTTP status code 400 Bad Request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

If you provide an individual checksum, Amazon S3 ignores any provided ChecksumAlgorithm parameter.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 | XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in XML format.

BucketLoggingStatus

Root level tag for the BucketLoggingStatus parameters.

Required: Yes

LoggingEnabled

Describes where logs are stored and the prefix that Amazon S3 assigns to all log object keys for a bucket. For more information, see [PUT Bucket logging](#) in the *Amazon S3 API Reference*.

Type: [LoggingEnabled](#) data type

Required: No

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Sample Request

This request enables logging and gives the grantee of the bucket READ access to the logs.

Buckets that use the bucket owner enforced setting for Object Ownership to disable ACLs don't support target grants. For more information, see [Permissions for server access log delivery](#) in the *Amazon S3 User Guide*.

```
PUT ?logging HTTP/1.1
Host: quotes.s3.<Region>.amazonaws.com
Content-Length: 214
Date: Wed, 25 Nov 2009 12:00:00 GMT
```

Authorization: authorization string

```
<?xml version="1.0" encoding="UTF-8"?>
<BucketLoggingStatus xmlns="http://doc.s3.amazonaws.com/2006-03-01">
  <LoggingEnabled>
    <TargetBucket>mybucketlogs</TargetBucket>
    <TargetPrefix>mybucket-access_log-</TargetPrefix>
    <TargetGrants>
      <Grant>
        <Grantee xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
          xsi:type="AmazonCustomerByEmail">
          <EmailAddress>user@company.com</EmailAddress>
        </Grantee>
        <Permission>READ</Permission>
      </Grant>
    </TargetGrants>
  </LoggingEnabled>
</BucketLoggingStatus>
```

Sample Response

This example illustrates one usage of PutBucketLogging.

```
HTTP/1.1 200 OK
x-amz-id-2: YgIPIfBiKa2bj0KMg95r/0zo3emzU4dzsD4rcKCHQUAdQkf3ShJT00pXUueF6QKo
x-amz-request-id: 236A8905248E5A01
Date: Wed, 01 Mar 2006 12:00:00 GMT
```

Sample Request: Disabling logging

This request disables logging on the bucket quotes.

```
PUT ?logging HTTP/1.1
Host: quotes.s3.<Region>.amazonaws.com
Content-Length: 214
Date: Wed, 25 Nov 2009 12:00:00 GMT
Authorization: authorization string

<?xml version="1.0" encoding="UTF-8"?>
```

```
<BucketLoggingStatus xmlns="http://doc.s3.amazonaws.com/2006-03-01" />
```

Sample Response

This example illustrates one usage of PutBucketLogging.

```
HTTP/1.1 200 OK
x-amz-id-2: YgIPIfBiKa2bj0KMg95r/0zo3emzU4dzsD4rcKCHQUAdQkf3ShJT00pXUueF6QKo
x-amz-request-id: 236A8905248E5A01
Date: Wed, 01 Mar 2006 12:00:00 GMT
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutBucketMetricsConfiguration

Service: Amazon S3

Sets a metrics configuration (specified by the metrics configuration ID) for the bucket. You can have up to 1,000 metrics configurations per bucket. If you're updating an existing metrics configuration, note that this is a full replacement of the existing metrics configuration. If you don't include the elements you want to keep, they are erased.

Note

Directory buckets - For directory buckets, you must make requests for this API operation to the Regional endpoint. These endpoints support path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name` . Virtual-hosted-style requests aren't supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

Permissions

To use this operation, you must have permissions to perform the `s3:PutMetricsConfiguration` action. The bucket owner has this permission by default. The bucket owner can grant this permission to others. For more information about permissions, see [Permissions Related to Bucket Subresource Operations](#) and [Managing Access Permissions to Your Amazon S3 Resources](#).

- **General purpose bucket permissions** - The `s3:PutMetricsConfiguration` permission is required in a policy. For more information about general purpose buckets permissions, see [Using Bucket Policies and User Policies](#) in the *Amazon S3 User Guide*.
- **Directory bucket permissions** - To grant access to this API operation, you must have the `s3express:PutMetricsConfiguration` permission in an IAM identity-based policy instead of a bucket policy. Cross-account access to this API operation isn't supported. This operation can only be performed by the AWS account that owns the resource. For more information about directory bucket policies and permissions, see [AWS Identity and Access Management \(IAM\) for S3 Express One Zone](#) in the *Amazon S3 User Guide*.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `s3express-control.region-code.amazonaws.com`.

For information about CloudWatch request metrics for Amazon S3, see [Monitoring Metrics with Amazon CloudWatch](#).

The following operations are related to `PutBucketMetricsConfiguration`:

- [DeleteBucketMetricsConfiguration](#)
- [GetBucketMetricsConfiguration](#)
- [ListBucketMetricsConfigurations](#)

`PutBucketMetricsConfiguration` has the following special error:

- Error code: `TooManyConfigurations`
 - Description: You are attempting to create a new configuration but have already reached the 1,000-configuration limit.
 - HTTP Status Code: HTTP 400 Bad Request

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /?metrics&id=Id HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<MetricsConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Id>string</Id>
  <Filter>
```

```

<AccessPointArn>string</AccessPointArn>
<And>
  <AccessPointArn>string</AccessPointArn>
  <Prefix>string</Prefix>
  <Tag>
    <Key>string</Key>
    <Value>string</Value>
  </Tag>
  ...
</And>
<Prefix>string</Prefix>
<Tag>
  <Key>string</Key>
  <Value>string</Value>
</Tag>
</Filter>
</MetricsConfiguration>

```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket for which the metrics configuration is set.

Directory buckets - When you use this operation with a directory bucket, you must use path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name`. Virtual-hosted-style requests aren't supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must also follow the format `bucket-base-name--zone-id--x-s3` (for example, `DOC-EXAMPLE-BUCKET--usw2-az1--x-s3`). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*.

Required: Yes

id

The ID used to identify the metrics configuration. The ID has a 64 character limit and can only contain letters, numbers, periods, dashes, and underscores.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Note

For directory buckets, this header is not supported in this API operation. If you specify this header, the request fails with the HTTP status code 501 Not Implemented.

Request Body

The request accepts the following data in XML format.

MetricsConfiguration

Root level tag for the MetricsConfiguration parameters.

Required: Yes

Filter

Specifies a metrics configuration filter. The metrics configuration will only include objects that meet the filter's criteria. A filter must be a prefix, an object tag, an access point ARN, or a conjunction (MetricsAndOperator).

Note

Metrics configurations for directory buckets do not support tag filters.

Type: [MetricsFilter](#) data type

Required: No

Id

The ID used to identify the metrics configuration. The ID has a 64 character limit and can only contain letters, numbers, periods, dashes, and underscores.

Type: String

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

First Sample Request

Put a metric configuration that enables metrics for an entire bucket.

```
PUT /?metrics&id=EntireBucket HTTP/1.1
Host: examplebucket.s3.<Region>.amazonaws.com
x-amz-date: Thu, 15 Nov 2016 00:17:21 GMT
Authorization: signatureValue
Content-Length: 159

<?xml version="1.0" encoding="UTF-8"?>
<MetricsConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Id>EntireBucket</Id>
</MetricsConfiguration>
```

First Sample Response

This example illustrates one usage of PutBucketMetricsConfiguration.

```
HTTP/1.1 204 No Content
x-amz-id-2: ITnGT1y4REXAMPLEPi4hklTXouTf0hccUjo0iCPEXAMPLEutBj3M7fPG1W02SEWp
x-amz-request-id: 51991EXAMPLE5321
Date: Thu, 15 Nov 2016 00:17:22 GMT
Server: AmazonS3
```

Second Sample Request

Put a metrics configuration that enables metrics for objects that start with a particular prefix and also have specific tags applied.

```
PUT /?metrics&id=ImportantBlueDocuments HTTP/1.1
Host: examplebucket.s3.<Region>.amazonaws.com
x-amz-date: Thu, 15 Nov 2016 00:17:29 GMT
Authorization: signatureValue
Content-Length: 480

<?xml version="1.0" encoding="UTF-8"?>
<MetricsConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Id>ImportantBlueDocuments</Id>
  <Filter>
    <And>
      <Prefix>documents/</Prefix>
      <Tag>
        <Key>priority</Key>
        <Value>high</Value>
      </Tag>
      <Tag>
        <Key>class</Key>
        <Value>blue</Value>
      </Tag>
    </And>
  </Filter>
</MetricsConfiguration>
```

Second Sample Response

This example illustrates one usage of PutBucketMetricsConfiguration.

```
HTTP/1.1 204 No Content
x-amz-id-2: ITnGT1y4REXAMPLEPi4hk1TXouTf0hccUjo0iCPEXAMPLEutBj3M7fPGLW02SEWp
x-amz-request-id: 51991EXAMPLE5321
Date: Thu, 15 Nov 2016 00:17:29 GMT
Server: AmazonS3
```

Third Sample Request

Put a metrics configuration that enables metrics for a specific access point.

```
PUT /?metrics&id=ImportantDocumentsAccessPoint HTTP/1.1
Host: examplebucket.s3.<Region>.amazonaws.com
x-amz-date: Thu, 26 Aug 2021 00:17:29 GMT
Authorization: signatureValue
Content-Length: 480

<?xml version="1.0" encoding="UTF-8"?>
<MetricsConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Id>ImportantDocumentsAccessPoint</Id>
  <Filter>
    <AccessPointArn>arn:aws:s3:us-west-2:123456789012:accesspoint/test</
AccessPointArn>
  </Filter>
</MetricsConfiguration>
```

Thirds Sample Response

This example illustrates one usage of PutBucketMetricsConfiguration.

```
HTTP/1.1 204 No Content
x-amz-id-2: ITnGT1y4REXAMPLEPi4hk1TXouTf0hccUjo0iCPEXAMPLEutBj3M7fPG1W02SEWp
x-amz-request-id: 51991EXAMPLE5321
Date: Thu, 26 Aug 2021 00:17:29 GMT
Server: AmazonS3
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)

- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutBucketNotification

Service: Amazon S3

Note

This operation is not supported for directory buckets.

No longer used, see the [PutBucketNotificationConfiguration](#) operation.

Request Syntax

```
PUT /?notification HTTP/1.1
Host: Bucket.s3.amazonaws.com
Content-MD5: ContentMD5
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<NotificationConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <TopicConfiguration>
    <Event>string</Event>
    <Event>string</Event>
    ...
    <Id>string</Id>
    <Topic>string</Topic>
  </TopicConfiguration>
  <QueueConfiguration>
    <Event>string</Event>
    <Event>string</Event>
    ...
    <Id>string</Id>
    <Queue>string</Queue>
  </QueueConfiguration>
  <CloudFunctionConfiguration>
    <CloudFunction>string</CloudFunction>
    <Event>string</Event>
    <Event>string</Event>
    ...
    <Id>string</Id>
    <InvocationRole>string</InvocationRole>
  </CloudFunctionConfiguration>
</NotificationConfiguration>
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket.

Required: Yes

Content-MD5

The MD5 hash of the PutPublicAccessBlock request body.

For requests made using the AWS Command Line Interface (CLI) or AWS SDKs, this field is calculated automatically.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-sdk-checksum-algorithm

Indicates the algorithm used to create the checksum for the request when you use the SDK. This header will not provide any additional functionality if you don't use the SDK. When you send this header, there must be a corresponding `x-amz-checksum` or `x-amz-trailer` header sent. Otherwise, Amazon S3 fails the request with the HTTP status code 400 Bad Request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

If you provide an individual checksum, Amazon S3 ignores any provided `ChecksumAlgorithm` parameter.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 | XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in XML format.

NotificationConfiguration

Root level tag for the NotificationConfiguration parameters.

Required: Yes

CloudFunctionConfiguration

Container for specifying the AWS Lambda notification configuration.

Type: [CloudFunctionConfiguration](#) data type

Required: No

QueueConfiguration

This data type is deprecated. This data type specifies the configuration for publishing messages to an Amazon Simple Queue Service (Amazon SQS) queue when Amazon S3 detects specified events.

Type: [QueueConfigurationDeprecated](#) data type

Required: No

TopicConfiguration

This data type is deprecated. A container for specifying the configuration for publication of messages to an Amazon Simple Notification Service (Amazon SNS) topic when Amazon S3 detects specified events.

Type: [TopicConfigurationDeprecated](#) data type

Required: No

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutBucketNotificationConfiguration

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Enables notifications of specified events for a bucket. For more information about event notifications, see [Configuring Event Notifications](#).

Using this API, you can replace an existing notification configuration. The configuration is an XML file that defines the event types that you want Amazon S3 to publish and the destination where you want Amazon S3 to publish an event notification when it detects an event of the specified type.

By default, your bucket has no event notifications configured. That is, the notification configuration will be an empty NotificationConfiguration.

```
<NotificationConfiguration>
```

```
</NotificationConfiguration>
```

This action replaces the existing notification configuration with the configuration you include in the request body.

After Amazon S3 receives this request, it first verifies that any Amazon Simple Notification Service (Amazon SNS) or Amazon Simple Queue Service (Amazon SQS) destination exists, and that the bucket owner has permission to publish to it by sending a test notification. In the case of AWS Lambda destinations, Amazon S3 verifies that the Lambda function permissions grant Amazon S3 permission to invoke the function from the Amazon S3 bucket. For more information, see [Configuring Notifications for Amazon S3 Events](#).

You can disable notifications by adding the empty NotificationConfiguration element.

For more information about the number of event notification configurations that you can create per bucket, see [Amazon S3 service quotas](#) in *AWS General Reference*.

By default, only the bucket owner can configure notifications on a bucket. However, bucket owners can use a bucket policy to grant permission to other users to set this configuration with the required `s3:PutBucketNotification` permission.

Note

The PUT notification is an atomic operation. For example, suppose your notification configuration includes SNS topic, SQS queue, and Lambda function configurations. When you send a PUT request with this configuration, Amazon S3 sends test messages to your SNS topic. If the message fails, the entire PUT action will fail, and Amazon S3 will not add the configuration to your bucket.

If the configuration in the request body includes only one `TopicConfiguration` specifying only the `s3:ReducedRedundancyLostObject` event type, the response will also include the `x-amz-sns-test-message-id` header containing the message ID of the test notification sent to the topic.

The following action is related to `PutBucketNotificationConfiguration`:

- [GetBucketNotificationConfiguration](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /?notification HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-expected-bucket-owner: ExpectedBucketOwner
x-amz-skip-destination-validation: SkipDestinationValidation
<?xml version="1.0" encoding="UTF-8"?>
<NotificationConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <TopicConfiguration>
    <Event>string</Event>
    ...
  <Filter>
    <S3Key>
      <FilterRule>
```

```

        <Name>string</Name>
        <Value>string</Value>
    </FilterRule>
    ...
</S3Key>
</Filter>
<Id>string</Id>
<Topic>string</Topic>
</TopicConfiguration>
...
<QueueConfiguration>
    <Event>string</Event>
    ...
    <Filter>
        <S3Key>
            <FilterRule>
                <Name>string</Name>
                <Value>string</Value>
            </FilterRule>
            ...
        </S3Key>
    </Filter>
    <Id>string</Id>
    <Queue>string</Queue>
</QueueConfiguration>
...
<CloudFunctionConfiguration>
    <Event>string</Event>
    ...
    <Filter>
        <S3Key>
            <FilterRule>
                <Name>string</Name>
                <Value>string</Value>
            </FilterRule>
            ...
        </S3Key>
    </Filter>
    <Id>string</Id>
    <CloudFunction>string</CloudFunction>
</CloudFunctionConfiguration>
...
<EventBridgeConfiguration>
</EventBridgeConfiguration>

```

```
</NotificationConfiguration>
```

URI Request Parameters

The request uses the following URI parameters.

[Bucket](#)

The name of the bucket.

Required: Yes

[x-amz-expected-bucket-owner](#)

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

[x-amz-skip-destination-validation](#)

Skips validation of Amazon SQS, Amazon SNS, and AWS Lambda destinations. True or false value.

Request Body

The request accepts the following data in XML format.

[NotificationConfiguration](#)

Root level tag for the NotificationConfiguration parameters.

Required: Yes

[CloudFunctionConfiguration](#)

Describes the AWS Lambda functions to invoke and the events for which to invoke them.

Type: Array of [LambdaFunctionConfiguration](#) data types

Required: No

[EventBridgeConfiguration](#)

Enables delivery of events to Amazon EventBridge.

Type: [EventBridgeConfiguration](#) data type

Required: No

[QueueConfiguration](#)

The Amazon Simple Queue Service queues to publish messages to and the events for which to publish messages.

Type: Array of [QueueConfiguration](#) data types

Required: No

[TopicConfiguration](#)

The topic to which notifications are sent and the events for which notifications are generated.

Type: Array of [TopicConfiguration](#) data types

Required: No

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Example 1: Configure notification to invoke a cloud function in Lambda

The following notification configuration includes `CloudFunctionConfiguration`, which identifies the event type for which Amazon S3 can invoke a cloud function and the name of the cloud function to invoke.

```
<NotificationConfiguration>
  <CloudFunctionConfiguration>
    <Id>ObjectCreatedEvents</Id>
    <CloudFunction>arn:aws:lambda:us-west-2:35667example:function:CreateThumbnail</
CloudFunction>
    <Event>s3:ObjectCreated:*</Event>
```

```
</CloudFunctionConfiguration>  
</NotificationConfiguration>
```

Example

The following PUT uploads the notification configuration. The action replaces the existing notification configuration.

```
PUT http://s3.<Region>.amazonaws.com/examplebucket?notification= HTTP/1.1  
User-Agent: s3curl 2.0  
Host: s3.amazonaws.com  
Pragma: no-cache  
Accept: */*  
Proxy-Connection: Keep-Alive  
Authorization: authorization string  
Date: Mon, 13 Oct 2014 23:14:52 +0000  
Content-Length: length  
  
[request body]
```

Sample Response

This example illustrates one usage of PutBucketNotificationConfiguration.

```
HTTP/1.1 200 OK  
x-amz-id-2: 8+FlwagBSoT2qpMaGlfcUkRkFR5W30eS7UhhoBb17j+kqvpS2cSF1gJ5coLd53d2  
x-amz-request-id: E5BA4600A3937335  
Date: Fri, 31 Oct 2014 01:49:50 GMT  
Content-Length: 0  
Server: AmazonS3
```

Example 2: Configure a notification with multiple destinations

The following notification configuration includes the topic and queue configurations:

- A topic configuration identifying an SNS topic for Amazon S3 to publish events of the `s3:ReducedRedundancyLostObject` type.

- A queue configuration identifying an SQS queue for Amazon S3 to publish events of the `s3:ObjectCreated:*` type.

```
<NotificationConfiguration>
  <TopicConfiguration>
    <Topic>arn:aws:sns:us-east-1:356671443308:s3notificationtopic2</Topic>
    <Event>s3:ReducedRedundancyLostObject</Event>
  </TopicConfiguration>
  <QueueConfiguration>
    <Queue>arn:aws:sqs:us-east-1:356671443308:s3notificationqueue</Queue>
    <Event>s3:ObjectCreated:*</Event>
  </QueueConfiguration>
</NotificationConfiguration>
```

Example

The following PUT request against the notification subresource of the `examplebucket` bucket sends the preceding notification configuration in the request body. The action replaces the existing notification configuration on the bucket.

```
PUT http://s3.<Region>.amazonaws.com/examplebucket?notification= HTTP/1.1
User-Agent: s3curl 2.0
Host: s3.amazonaws.com
Pragma: no-cache
Accept: */*
Proxy-Connection: Keep-Alive
Authorization: authorization string
Date: Mon, 13 Oct 2014 22:58:43 +0000
Content-Length: 391
Expect: 100-continue
```

Example 3: Configure a notification with object key name filtering

The following notification configuration contains a queue configuration identifying an Amazon SQS queue for Amazon S3 to publish events to of the `s3:ObjectCreated:Put` type. The events will be published whenever an object that has a prefix of `images/` and a `.jpg` suffix is PUT to a

bucket. For more examples of notification configurations that use filtering, see [Configuring Event Notifications](#).

```
<NotificationConfiguration>
  <QueueConfiguration>
    <Id>1</Id>
    <Filter>
      <S3Key>
        <FilterRule>
          <Name>prefix</Name>
          <Value>images/</Value>
        </FilterRule>
        <FilterRule>
          <Name>suffix</Name>
          <Value>.jpg</Value>
        </FilterRule>
      </S3Key>
    </Filter>
    <Queue>arn:aws:sqs:us-west-2:444455556666:s3notificationqueue</Queue>
    <Event>s3:ObjectCreated:Put</Event>
  </QueueConfiguration>
</NotificationConfiguration>
```

Example

The following PUT request against the notification subresource of the `examplebucket` bucket sends the preceding notification configuration in the request body. The action replaces the existing notification configuration on the bucket.

```
PUT http://s3.<Region>.amazonaws.com/examplebucket?notification= HTTP/1.1
User-Agent: s3curl 2.0
Host: s3.amazonaws.com
Pragma: no-cache
Accept: */*
Proxy-Connection: Keep-Alive
Authorization: authorization string
Date: Mon, 13 Oct 2014 22:58:43 +0000
Content-Length: length
Expect: 100-continue
```

Sample Response

This example illustrates one usage of PutBucketNotificationConfiguration.

```
HTTP/1.1 200 OK
x-amz-id-2: SlvJLkfunoAGILZK3KqHSSUq4kwbudk1R0mESoH0pDacULy+cxRoR1Sv1foyvvg2A
x-amz-request-id: BB1BA8E12D6A80B7
Date: Mon, 13 Oct 2014 22:58:44 GMT
Content-Length: 0
Server: AmazonS3
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutBucketOwnershipControls

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Creates or modifies OwnershipControls for an Amazon S3 bucket. To use this operation, you must have the `s3:PutBucketOwnershipControls` permission. For more information about Amazon S3 permissions, see [Specifying permissions in a policy](#).

For information about Amazon S3 Object Ownership, see [Using object ownership](#).

The following operations are related to PutBucketOwnershipControls:

- [GetBucketOwnershipControls](#)
- [DeleteBucketOwnershipControls](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /?ownershipControls HTTP/1.1
Host: Bucket.s3.amazonaws.com
Content-MD5: ContentMD5
x-amz-expected-bucket-owner: ExpectedBucketOwner
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
<?xml version="1.0" encoding="UTF-8"?>
<OwnershipControls xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Rule>
    <ObjectOwnership>string</ObjectOwnership>
  </Rule>
  ...
```

```
</OwnershipControls>
```

URI Request Parameters

The request uses the following URI parameters.

[Bucket](#)

The name of the Amazon S3 bucket whose `OwnershipControls` you want to set.

Required: Yes

[Content-MD5](#)

The MD5 hash of the `OwnershipControls` request body.

For requests made using the AWS Command Line Interface (CLI) or AWS SDKs, this field is calculated automatically.

[x-amz-expected-bucket-owner](#)

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code `403 Forbidden` (access denied).

[x-amz-sdk-checksum-algorithm](#)

Indicates the algorithm used to create the checksum for the object when you use the SDK. This header will not provide any additional functionality if you don't use the SDK. When you send this header, there must be a corresponding `x-amz-checksum-algorithm` header sent. Otherwise, Amazon S3 fails the request with the HTTP status code `400 Bad Request`. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

If you provide an individual checksum, Amazon S3 ignores any provided `ChecksumAlgorithm` parameter.

Valid Values: `CRC32` | `CRC32C` | `SHA1` | `SHA256` | `CRC64NVME` | `SHA512` | `MD5` | `XXHASH64` | `XXHASH3` | `XXHASH128`

Request Body

The request accepts the following data in XML format.

OwnershipControls

Root level tag for the OwnershipControls parameters.

Required: Yes

Rule

The container element for an ownership control rule.

Type: Array of [OwnershipControlsRule](#) data types

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Sample Request with BucketOwnerEnforced OwnershipControls

The following request puts a bucket OwnershipControls that specifies BucketOwnerEnforced.

```
PUT /amzn-s3-demo-bucket?ownershipControls= HTTP/1.1
Host:amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
x-amz-date: 20211130T230132Z
x-amz-content-sha256:
bafb46c18574a73704c8227aef060df1c12ea0d964e19b949d06e9f763805fe2
Authorization: authorization string

<?xml version="1.0" encoding="UTF-8"?>
<OwnershipControls xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Rule>
    <ObjectOwnership>BucketOwnerEnforced</ObjectOwnership>
  </Rule>
```



```
</OwnershipControls>
```

Sample Response with BucketOwnerEnforced OwnershipControls

This example illustrates one usage of PutBucketOwnershipControls.

```
HTTP/1.1 200 OK
x-amz-id-2: zkDVX0gbz8oKcjNz7GPz8XhXkhNArHtA8/
W0f5hyEj6SbisSRdqITZvSuAMik7HK4PY+izDZZI0=
x-amz-request-id: BK7Y8M3G7Z0RFRCF
Date: Tue, 30 Nov 2021 23:01:33 GMT
Content-Length: 0
Server: AmazonS3
```

Sample Request with BucketOwnerPreferred OwnershipControls

The following request puts a bucket OwnershipControls that specifies BucketOwnerPreferred.

```
PUT /amzn-s3-demo-bucket?ownershipControls= HTTP/1.1
Host:amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
x-amz-date: 20200618T230132Z
x-amz-content-sha256:
bafb46c18574a73704c8227aef060df1c12ea0d964e19b949d06e9f763805fe2
Authorization: authorization string

<?xml version="1.0" encoding="UTF-8"?>
<OwnershipControls xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Rule>
    <ObjectOwnership>BucketOwnerPreferred</ObjectOwnership>
  </Rule>
</OwnershipControls>
```

Sample Response with BucketOwnerPreferred OwnershipControls

This example illustrates one usage of PutBucketOwnershipControls.

```
HTTP/1.1 200 OK
x-amz-id-2: zkDVX0gbz8oKcjNz7GPz8XhXkhNArHtA8/
W0f5hyEj6SbisSRdqITZvSuAMik7HK4PY+izDZZI0=
x-amz-request-id: BK7Y8M3G7Z0RFRCP
Date: Thu, 18 Jun 2020 23:01:33 GMT
Content-Length: 0
Server: AmazonS3
```

Sample Request with ObjectWriter OwnershipControls

The following request puts a bucket OwnershipControls that specifies ObjectWriter.

```
PUT /amzn-s3-demo-bucket?ownershipControls= HTTP/1.1
Host:amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
x-amz-date: 20200618T230132Z
x-amz-content-sha256:
bafb46c18574a73704c8227aef060df1c12ea0d964e19b949d06e9f763805fe2
Authorization: authorization string

<?xml version="1.0" encoding="UTF-8"?>
<OwnershipControls xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Rule>
    <ObjectOwnership>ObjectWriter</ObjectOwnership>
  </Rule>
</OwnershipControls>
```

Sample Response with ObjectWriter OwnershipControls

This example illustrates one usage of PutBucketOwnershipControls.

```
HTTP/1.1 200 OK
x-amz-id-2: zkDVX0gbz8oKcjNz7GPz8XhXkhNArHtA8/
W0f5hyEj6SbisSRdqITZvSuAMik7HK4PY+izDZZI0=
x-amz-request-id: BK7Y8M3G7Z0RFRCP
Date: Thu, 18 Jun 2020 23:01:33 GMT
Content-Length: 0
Server: AmazonS3
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutBucketPolicy

Service: Amazon S3

Applies an Amazon S3 bucket policy to an Amazon S3 bucket.

Note

Directory buckets - For directory buckets, you must make requests for this API operation to the Regional endpoint. These endpoints support path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name` . Virtual-hosted-style requests aren't supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

Permissions

If you are using an identity other than the root user of the AWS account that owns the bucket, the calling identity must both have the PutBucketPolicy permissions on the specified bucket and belong to the bucket owner's account in order to use this operation.

If you don't have PutBucketPolicy permissions, Amazon S3 returns a 403 Access Denied error. If you have the correct permissions, but you're not using an identity that belongs to the bucket owner's account, Amazon S3 returns a 405 Method Not Allowed error.

Important

To ensure that bucket owners don't inadvertently lock themselves out of their own buckets, the root principal in a bucket owner's AWS account can perform the GetBucketPolicy, PutBucketPolicy, and DeleteBucketPolicy API actions, even if their bucket policy explicitly denies the root principal's access. Bucket owner root principals can only be blocked from performing these API actions by VPC endpoint policies and AWS Organizations policies.

- **General purpose bucket permissions** - The `s3:PutBucketPolicy` permission is required in a policy. For more information about general purpose buckets bucket policies, see [Using Bucket Policies and User Policies](#) in the *Amazon S3 User Guide*.

- **Directory bucket permissions** - To grant access to this API operation, you must have the `s3express:PutBucketPolicy` permission in an IAM identity-based policy instead of a bucket policy. Cross-account access to this API operation isn't supported. This operation can only be performed by the AWS account that owns the resource. For more information about directory bucket policies and permissions, see [AWS Identity and Access Management \(IAM\) for S3 Express One Zone](#) in the *Amazon S3 User Guide*.

Example bucket policies

General purpose buckets example bucket policies - See [Bucket policy examples](#) in the *Amazon S3 User Guide*.

Directory bucket example bucket policies - See [Example bucket policies for S3 Express One Zone](#) in the *Amazon S3 User Guide*.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `s3express-control.region-code.amazonaws.com`.

The following operations are related to `PutBucketPolicy`:

- [CreateBucket](#)
- [DeleteBucket](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /?policy HTTP/1.1
Host: Bucket.s3.amazonaws.com
Content-MD5: ContentMD5
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-confirm-remove-self-bucket-access: ConfirmRemoveSelfBucketAccess
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

```
{ Policy in JSON format }
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket.

Directory buckets - When you use this operation with a directory bucket, you must use path-style requests in the format `https://s3express-control.region-code.amazonaws.com/bucket-name`. Virtual-hosted-style requests aren't supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must also follow the format `bucket-base-name--zone-id--x-s3` (for example, `DOC-EXAMPLE-BUCKET--usw2-az1--x-s3`). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*

Required: Yes

Content-MD5

The MD5 hash of the request body.

For requests made using the AWS Command Line Interface (CLI) or AWS SDKs, this field is calculated automatically.

Note

This functionality is not supported for directory buckets.

x-amz-confirm-remove-self-bucket-access

Set this parameter to true to confirm that you want to remove your permissions to change this bucket policy in the future.

Note

This functionality is not supported for directory buckets.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Note

For directory buckets, this header is not supported in this API operation. If you specify this header, the request fails with the HTTP status code 501 Not Implemented.

x-amz-sdk-checksum-algorithm

Indicates the algorithm used to create the checksum for the request when you use the SDK. This header will not provide any additional functionality if you don't use the SDK. When you send this header, there must be a corresponding `x-amz-checksum-algorithm` or `x-amz-trailer` header sent. Otherwise, Amazon S3 fails the request with the HTTP status code 400 Bad Request.

For the `x-amz-checksum-algorithm` header, replace *algorithm* with the supported algorithm from the following list:

- CRC32
- CRC32C
- CRC64NVME
- MD5
- SHA1
- SHA256
- SHA512
- XXHASH3
- XXHASH64
- XXHASH128

For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

If the individual checksum value you provide through `x-amz-checksum-algorithm` doesn't match the checksum algorithm you set through `x-amz-sdk-checksum-algorithm`, Amazon S3 fails the request with a `BadDigest` error.

 Note

For directory buckets, when you use AWS SDKs, CRC32 is the default checksum algorithm that's used for performance.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 | XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in JSON format.

[Policy](#)

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Sample Request for general purpose buckets

The following request shows the PUT individual policy request for the bucket.

```
PUT /?policy HTTP/1.1
Host: bucket.s3.<Region>.amazonaws.com
Date: Tue, 04 Apr 2010 20:34:56 GMT
Authorization: authorization string
```



```
{
  "Version": "2008-10-17",
  "Id": "aaaa-bbbb-cccc-dddd",
  "Statement" : [
    {
      "Effect": "Allow",
      "Sid": "1",
      "Principal" : {
        "AWS": ["111122223333", "444455556666"]
      },
      "Action": ["s3:*"],
      "Resource": "arn:aws:s3:::bucket/*"
    }
  ]
}
```

Sample Response for general purpose buckets

This example illustrates one usage of PutBucketPolicy.

```
HTTP/1.1 204 No Content
x-amz-id-2: Uuag1LuByR50nimru9SAMPLEAtRPfTa0Fg==
x-amz-request-id: 656c76696e6727732SAMPLE7374
Date: Tue, 04 Apr 2010 20:34:56 GMT
Connection: keep-alive
Server: AmazonS3
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)

- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutBucketReplication

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Creates a replication configuration or replaces an existing one. For more information, see [Replication](#) in the *Amazon S3 User Guide*.

Specify the replication configuration in the request body. In the replication configuration, you provide the name of the destination bucket or buckets where you want Amazon S3 to replicate objects, the IAM role that Amazon S3 can assume to replicate objects on your behalf, and other relevant information. You can invoke this request for a specific AWS Region by using the [aws:RequestedRegion](#) condition key.

A replication configuration must include at least one rule, and can contain a maximum of 1,000. Each rule identifies a subset of objects to replicate by filtering the objects in the source bucket. To choose additional subsets of objects to replicate, add a rule for each subset.

To specify a subset of the objects in the source bucket to apply a replication rule to, add the Filter element as a child of the Rule element. You can filter objects based on an object key prefix, one or more object tags, or both. When you add the Filter element in the configuration, you must also add the following elements: DeleteMarkerReplication, Status, and Priority.

Note

If you are using an earlier version of the replication configuration, Amazon S3 handles replication of delete markers differently. For more information, see [Backward Compatibility](#).

For information about enabling versioning on a bucket, see [Using Versioning](#).

Handling Replication of Encrypted Objects

By default, Amazon S3 doesn't replicate objects that are stored at rest using server-side encryption with KMS keys. To replicate AWS KMS-encrypted objects, add the

following: `SourceSelectionCriteria`, `SseKmsEncryptedObjects`, `Status`, `EncryptionConfiguration`, and `ReplicaKmsKeyID`. For information about replication configuration, see [Replicating Objects Created with SSE Using KMS keys](#).

For information on `PutBucketReplication` errors, see [List of replication-related error codes](#)

Permissions

To create a `PutBucketReplication` request, you must have `s3:PutReplicationConfiguration` permissions for the bucket.

By default, a resource owner, in this case the AWS account that created the bucket, can perform this operation. The resource owner can also grant others permissions to perform the operation. For more information about permissions, see [Specifying Permissions in a Policy](#) and [Managing Access Permissions to Your Amazon S3 Resources](#).

Note

To perform this operation, the user or role performing the action must have the [iam:PassRole](#) permission.

The following operations are related to `PutBucketReplication`:

- [GetBucketReplication](#)
- [DeleteBucketReplication](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /?replication HTTP/1.1
Host: Bucket.s3.amazonaws.com
Content-MD5: ContentMD5
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
```

```

x-amz-bucket-object-lock-token: Token
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<ReplicationConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Role>string</Role>
  <Rule>
    <DeleteMarkerReplication>
      <Status>string</Status>
    </DeleteMarkerReplication>
    <Destination>
      <AccessControlTranslation>
        <Owner>string</Owner>
      </AccessControlTranslation>
      <Account>string</Account>
      <Bucket>string</Bucket>
      <EncryptionConfiguration>
        <ReplicaKmsKeyID>string</ReplicaKmsKeyID>
      </EncryptionConfiguration>
      <Metrics>
        <EventThreshold>
          <Minutes>integer</Minutes>
        </EventThreshold>
        <Status>string</Status>
      </Metrics>
      <ReplicationTime>
        <Status>string</Status>
        <Time>
          <Minutes>integer</Minutes>
        </Time>
      </ReplicationTime>
      <StorageClass>string</StorageClass>
    </Destination>
    <ExistingObjectReplication>
      <Status>string</Status>
    </ExistingObjectReplication>
    <Filter>
      <And>
        <Prefix>string</Prefix>
        <Tag>
          <Key>string</Key>
          <Value>string</Value>
        </Tag>
        ...
      </And>

```

```

    <Prefix>string</Prefix>
    <Tag>
      <Key>string</Key>
      <Value>string</Value>
    </Tag>
  </Filter>
  <ID>string</ID>
  <Prefix>string</Prefix>
  <Priority>integer</Priority>
  <SourceSelectionCriteria>
    <ReplicaModifications>
      <Status>string</Status>
    </ReplicaModifications>
    <SseKmsEncryptedObjects>
      <Status>string</Status>
    </SseKmsEncryptedObjects>
  </SourceSelectionCriteria>
  <Status>string</Status>
</Rule>
...
</ReplicationConfiguration>

```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket

Required: Yes

Content-MD5

The Base64 encoded 128-bit MD5 digest of the data. You must use this header as a message integrity check to verify that the request body was not corrupted in transit. For more information, see [RFC 1864](#).

For requests made using the AWS Command Line Interface (CLI) or AWS SDKs, this field is calculated automatically.

x-amz-bucket-object-lock-token

A token to allow Object Lock to be enabled for an existing bucket.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-sdk-checksum-algorithm

Indicates the algorithm used to create the checksum for the request when you use the SDK. This header will not provide any additional functionality if you don't use the SDK. When you send this header, there must be a corresponding x-amz-checksum or x-amz-trailer header sent. Otherwise, Amazon S3 fails the request with the HTTP status code 400 Bad Request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

If you provide an individual checksum, Amazon S3 ignores any provided ChecksumAlgorithm parameter.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 | XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in XML format.

ReplicationConfiguration

Root level tag for the ReplicationConfiguration parameters.

Required: Yes

Role

The Amazon Resource Name (ARN) of the AWS Identity and Access Management (IAM) role that Amazon S3 assumes when replicating objects. For more information, see [How to Set Up Replication](#) in the *Amazon S3 User Guide*.

Type: String

Required: Yes

Rule

A container for one or more replication rules. A replication configuration must have at least one rule and can contain a maximum of 1,000 rules.

Type: Array of [ReplicationRule](#) data types

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Sample Request: Add a replication configuration

The following is a sample PUT request that creates a replication subresource on the specified bucket and saves the replication configuration in it. The replication configuration specifies a rule to replicate objects to the DOC-EXAMPLE-BUCKET bucket. The rule includes a filter to replicate only the objects created with the key name prefix TaxDocs and that have two specific tags.

After you add a replication configuration to your bucket, Amazon S3 assumes the AWS Identity and Access Management (IAM) role specified in the configuration to replicate objects on behalf of the bucket owner. The bucket owner is the AWS account that created the bucket.

Filtering using the <Filter> element is supported in the latest XML configuration. If you are using an earlier version of the XML configuration, you can filter only on key prefix. In that case, you add the <Prefix> element as a child of the <Rule>.

For more examples of replication configuration, see [Replication Configuration Overview](#) in the *Amazon S3 User Guide*.

```
PUT /?replication HTTP/1.1
Host: examplebucket.s3.<Region>.amazonaws.com
Date: Wed, 11 Feb 2015 02:11:21 GMT
Content-MD5: q6yJDlIkBaGGfb3QLY69A==
Authorization: authorization string
Content-Length: length
```



```
<ReplicationConfiguration>
  <Role>arn:aws:iam::35667example:role/CrossRegionReplicationRoleForS3</Role>
  <Rule>
    <ID>rule1</ID>
    <Status>Enabled</Status>
    <Priority>1</Priority>
    <DeleteMarkerReplication>
      <Status>Disabled</Status>
    </DeleteMarkerReplication>
    <Filter>
      <And>
        <Prefix>TaxDocs</Prefix>
        <Tag>
          <Key>key1</Key>
          <Value>value1</Value>
        </Tag>
        <Tag>
          <Key>key1</Key>
          <Value>value1</Value>
        </Tag>
      </And>
    </Filter>
    <Destination>
      <Bucket>arn:aws:s3::DOC-EXAMPLE-BUCKET</Bucket>
    </Destination>
  </Rule>
</ReplicationConfiguration>
```

Sample Response

This example illustrates one usage of PutBucketReplication.

```
HTTP/1.1 200 OK
x-amz-id-2: r+qR7+nhXtJDDIJ0JJYcd+1j5nM/rUFiiiZ/fNbD0sd3JUE8NWMLNHXmvPfwMpdC
x-amz-request-id: 9E26D08072A8EF9E
Date: Wed, 11 Feb 2015 02:11:22 GMT
Content-Length: 0
Server: AmazonS3
```

Sample Request: Add a Replication Configuration with Amazon S3 Replication Time Control Enabled

You can use S3 Replication Time Control (S3 RTC) to replicate your data in the same AWS Region or across different AWS Regions in a predictable time frame. S3 RTC replicates 99.99 percent of new objects stored in Amazon S3 within 15 minutes. For more information, see [Replicating objects using Replication Time Control](#).

```
PUT /?replication HTTP/1.1
Host: examplebucket.s3.<Region>.amazonaws.com
Date: Wed, 11 Feb 2015 02:11:21 GMT
Content-MD5: q6yJDlIkBaGGfb3QLY69A==
Authorization: authorization string
Content-Length: length
x-amz-bucket-object-lock-token: Token
<?xml version="1.0" encoding="UTF-8"?>

<ReplicationConfiguration>
  <Role>arn:aws:iam::35667example:role/CrossRegionReplicationRoleForS3</Role>
  <Rule>
    <ID>rule1</ID>
    <Status>Enabled</Status>
    <Priority>1</Priority>
    <Filter>
      <And>
        <Prefix>TaxDocs</Prefix>
        <Tag>
          <Key>key1</Key>
          <Value>value1</Value>
        </Tag>
        <Tag>
          <Key>key1</Key>
          <Value>value1</Value>
        </Tag>
      </And>
    </Filter>
    <Destination>
      <Bucket>arn:aws:s3:::DOC-EXAMPLE-BUCKET</Bucket>
      <Metrics>
        <Status>Enabled</Status>
        <EventThreshold>
          <Minutes>15</Minutes>
```

```
        </EventThreshold>
    </Metrics>
    <ReplicationTime>
        <Status>Enabled</Status>
        <Time>
            <Minutes>15</Minutes>
        </Time>
    </ReplicationTime>
</Destination>
</Rule>
</ReplicationConfiguration>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutBucketRequestPayment

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Sets the request payment configuration for a bucket. By default, the bucket owner pays for downloads from the bucket. This configuration parameter enables the bucket owner (only) to specify that the person requesting the download will be charged for the download. For more information, see [Requester Pays Buckets](#).

The following operations are related to PutBucketRequestPayment:

- [CreateBucket](#)
- [GetBucketRequestPayment](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /?requestPayment HTTP/1.1
Host: Bucket.s3.amazonaws.com
Content-MD5: ContentMD5
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<RequestPaymentConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Payer>string</Payer>
</RequestPaymentConfiguration>
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name.

Required: Yes

Content-MD5

The Base64 encoded 128-bit MD5 digest of the data. You must use this header as a message integrity check to verify that the request body was not corrupted in transit. For more information, see [RFC 1864](#).

For requests made using the AWS Command Line Interface (CLI) or AWS SDKs, this field is calculated automatically.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-sdk-checksum-algorithm

Indicates the algorithm used to create the checksum for the request when you use the SDK. This header will not provide any additional functionality if you don't use the SDK. When you send this header, there must be a corresponding `x-amz-checksum` or `x-amz-trailer` header sent. Otherwise, Amazon S3 fails the request with the HTTP status code 400 Bad Request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

If you provide an individual checksum, Amazon S3 ignores any provided `ChecksumAlgorithm` parameter.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 | XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in XML format.

RequestPaymentConfiguration

Root level tag for the RequestPaymentConfiguration parameters.

Required: Yes

Payer

Specifies who pays for the download and request fees.

Type: String

Valid Values: Requester | BucketOwner

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Sample Request

This request creates a Requester Pays bucket named colorpictures.

```
PUT ?requestPayment HTTP/1.1
Host: colorpictures.s3.<Region>.amazonaws.com
Content-Length: 173
Date: Wed, 01 Mar 2006 12:00:00 GMT
Authorization: authorization string

<RequestPaymentConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Payer>Requester</Payer>
</RequestPaymentConfiguration>
```

Sample Response

Delete the metric configuration with a specified ID, which disables the CloudWatch metrics with the `ExampleMetrics` value for the `FilterId` dimension.

```
HTTP/1.1 200 OK
x-amz-id-2: YgIPIfBiKa2bj0KMg95r/0zo3emzU4dzsD4rcKCHQUAdQkf3ShJT00pXUueF6QKo
x-amz-request-id: 236A8905248E5A01
Date: Wed, 01 Mar 2006 12:00:00 GMT
Location: /colorpictures
Content-Length: 0
Connection: close
Server: AmazonS3
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutBucketTagging

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Sets the tags for a general purpose bucket if attribute based access control (ABAC) is not enabled for the bucket. When you [enable ABAC for a general purpose bucket](#), you can no longer use this operation for that bucket and must use the [TagResource](#) or [UntagResource](#) operations instead.

Use tags to organize your AWS bill to reflect your own cost structure. To do this, sign up to get your AWS account bill with tag key values included. Then, to see the cost of combined resources, organize your billing information according to resources with the same tag key values. For example, you can tag several resources with a specific application name, and then organize your billing information to see the total cost of that application across several services. For more information, see [Cost Allocation and Tagging](#) and [Using Cost Allocation in Amazon S3 Bucket Tags](#).

Note

When this operation sets the tags for a bucket, it will overwrite any current tags the bucket already has. You cannot use this operation to add tags to an existing list of tags.

To use this operation, you must have permissions to perform the `s3:PutBucketTagging` action. The bucket owner has this permission by default and can grant this permission to others. For more information about permissions, see [Permissions Related to Bucket Subresource Operations](#) and [Managing Access Permissions to Your Amazon S3 Resources](#).

PutBucketTagging has the following special errors. For more Amazon S3 errors see, [Error Responses](#).

- **InvalidTag** - The tag provided was not a valid tag. This error can occur if the tag did not pass input validation. For more information, see [Using Cost Allocation in Amazon S3 Bucket Tags](#).
- **MalformedXML** - The XML provided does not match the schema.
- **OperationAborted** - A conflicting conditional action is currently in progress against this resource. Please try again.

- `InternalServerError` - The service was unable to apply the provided tag to the bucket.

The following operations are related to `PutBucketTagging`:

- [GetBucketTagging](#)
- [DeleteBucketTagging](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /?tagging HTTP/1.1
Host: Bucket.s3.amazonaws.com
Content-MD5: ContentMD5
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<Tagging xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <TagSet>
    <Tag>
      <Key>string</Key>
      <Value>string</Value>
    </Tag>
  </TagSet>
</Tagging>
```

URI Request Parameters

The request uses the following URI parameters.

[Bucket](#)

The bucket name.

Required: Yes

Content-MD5

The Base64 encoded 128-bit MD5 digest of the data. You must use this header as a message integrity check to verify that the request body was not corrupted in transit. For more information, see [RFC 1864](#).

For requests made using the AWS Command Line Interface (CLI) or AWS SDKs, this field is calculated automatically.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-sdk-checksum-algorithm

Indicates the algorithm used to create the checksum for the request when you use the SDK. This header will not provide any additional functionality if you don't use the SDK. When you send this header, there must be a corresponding `x-amz-checksum` or `x-amz-trailer` header sent. Otherwise, Amazon S3 fails the request with the HTTP status code 400 Bad Request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

If you provide an individual checksum, Amazon S3 ignores any provided `ChecksumAlgorithm` parameter.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 | XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in XML format.

Tagging

Root level tag for the Tagging parameters.

Required: Yes

TagSet

A collection for a set of tags

Type: Array of [Tag](#) data types

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Sample Request: Add tag set to a bucket

The following request adds a tag set to the existing `examplebucket` bucket.

```
PUT ?tagging HTTP/1.1
Host: examplebucket.s3.<Region>.amazonaws.com
Content-Length: 1660
x-amz-date: Thu, 12 Apr 2012 20:04:21 GMT
Authorization: authorization string
```

```
<Tagging>
  <TagSet>
    <Tag>
      <Key>Project</Key>
      <Value>Project One</Value>
    </Tag>
    <Tag>
      <Key>User</Key>
      <Value>jsmith</Value>
    </Tag>
  </TagSet>
</Tagging>
```

Sample Response

This example illustrates one usage of `PutBucketTagging`.

```
HTTP/1.1 204 No Content
x-amz-id-2: YgIPIfBiKa2bj0KMgUAdQkf3ShJT00pXUueF6QKo
x-amz-request-id: 236A8905248E5A01
Date: Wed, 01 Oct 2012 12:00:00 GMT
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutBucketVersioning

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Note

When you enable versioning on a bucket for the first time, it might take a short amount of time for the change to be fully propagated. While this change is propagating, you might encounter intermittent HTTP 404 NoSuchKey errors for requests to objects created or updated after enabling versioning. We recommend that you wait for 15 minutes after enabling versioning before issuing write operations (PUT or DELETE) on objects in the bucket.

Sets the versioning state of an existing bucket.

You can set the versioning state with one of the following values:

Enabled—Enables versioning for the objects in the bucket. All objects added to the bucket receive a unique version ID.

Suspended—Disables versioning for the objects in the bucket. All objects added to the bucket receive the version ID null.

If the versioning state has never been set on a bucket, it has no versioning state; a [GetBucketVersioning](#) request does not return a versioning state value.

In order to enable MFA Delete, you must be the bucket owner. If you are the bucket owner and want to enable MFA Delete in the bucket versioning configuration, you must include the `x-amz-mfa` request header and the `Status` and the `MfaDelete` request elements in a request to set the versioning state of the bucket.

Important

If you have an object expiration lifecycle configuration in your non-versioned bucket and you want to maintain the same permanent delete behavior when you enable

versioning, you must add a noncurrent expiration policy. The noncurrent expiration lifecycle configuration will manage the deletes of the noncurrent object versions in the version-enabled bucket. (A version-enabled bucket maintains one current and zero or more noncurrent object versions.) For more information, see [Lifecycle and Versioning](#).

The following operations are related to PutBucketVersioning:

- [CreateBucket](#)
- [DeleteBucket](#)
- [GetBucketVersioning](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /?versioning HTTP/1.1
Host: Bucket.s3.amazonaws.com
Content-MD5: ContentMD5
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-mfa: MFA
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<VersioningConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <MfaDelete>string</MfaDelete>
  <Status>string</Status>
</VersioningConfiguration>
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name.

Required: Yes

Content-MD5

>The Base64 encoded 128-bit MD5 digest of the data. You must use this header as a message integrity check to verify that the request body was not corrupted in transit. For more information, see [RFC 1864](#).

For requests made using the AWS Command Line Interface (CLI) or AWS SDKs, this field is calculated automatically.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-mfa

The concatenation of the authentication device's serial number, a space, and the value that is displayed on your authentication device. The serial number is the number that uniquely identifies the MFA device. For physical MFA devices, this is the unique serial number that's provided with the device. For virtual MFA devices, the serial number is the device ARN. For more information, see [Enabling versioning on buckets](#) and [Configuring MFA delete](#) in the *Amazon Simple Storage Service User Guide*.

x-amz-sdk-checksum-algorithm

Indicates the algorithm used to create the checksum for the request when you use the SDK. This header will not provide any additional functionality if you don't use the SDK. When you send this header, there must be a corresponding `x-amz-checksum` or `x-amz-trailer` header sent. Otherwise, Amazon S3 fails the request with the HTTP status code 400 Bad Request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

If you provide an individual checksum, Amazon S3 ignores any provided `ChecksumAlgorithm` parameter.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 | XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in XML format.

VersioningConfiguration

Root level tag for the VersioningConfiguration parameters.

Required: Yes

MFADelete

Specifies whether MFA delete is enabled in the bucket versioning configuration. This element is only returned if the bucket has been configured with MFA delete. If the bucket has never been so configured, this element is not returned.

Type: String

Valid Values: Enabled | Disabled

Required: No

Status

The versioning state of the bucket.

Type: String

Valid Values: Enabled | Suspended

Required: No

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Sample Request

The following request enables versioning for the specified bucket.


```
PUT /?versioning HTTP/1.1
Host: bucket.s3.<Region>.amazonaws.com
Date: Wed, 01 Mar 2006 12:00:00 GMT
Authorization: authorization string
Content-Type: text/plain
Content-Length: 124

<VersioningConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Status>Enabled</Status>
</VersioningConfiguration>
```

Sample Response

This example illustrates one usage of PutBucketVersioning.

```
HTTP/1.1 200 OK
x-amz-id-2: YgIPIfBiKa2bj0KMg95r/0zo3emzU4dzsD4rcKCHQUAdQkf3ShJT00pXUueF6QKo
x-amz-request-id: 236A8905248E5A01
Date: Wed, 01 Mar 2006 12:00:00 GMT3
```

Sample Request

The following request suspends versioning for the specified bucket.

```
PUT /?versioning HTTP/1.1
Host: bucket.s3.<Region>.amazonaws.com
Date: Wed, 12 Oct 2009 17:50:00 GMT
Authorization: authorization string
Content-Type: text/plain
Content-Length: 124

<VersioningConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Status>Suspended</Status>
</VersioningConfiguration>
```

Sample Response

This example illustrates one usage of PutBucketVersioning.

```
HTTP/1.1 200 OK
x-amz-id-2: YgIPIfBiKa2bj0KMg95r/0zo3emzU4dzsD4rcKCHQUAdQkf3ShJT00pXUueF6QKo
x-amz-request-id: 236A8905248E5A01
Date: Wed, 01 Mar 2006 12:00:00 GMT
```

Sample Request

The following request enables versioning and MFA Delete on a bucket. Note the space between [SerialNumber] and [TokenCode] and that you must include Status whenever you use MfaDelete.

```
PUT /?versioning HTTP/1.1
Host: bucket.s3.<Region>.amazonaws.com
Date: Wed, 12 Oct 2009 17:50:00 GMT
x-amz-mfa:[SerialNumber] [TokenCode]
Authorization: authorization string
Content-Type: text/plain
Content-Length: 124

<VersioningConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Status>Enabled</Status>
  <MfaDelete>Enabled</MfaDelete>
</VersioningConfiguration>
```

Sample Response

This example illustrates one usage of PutBucketVersioning.

```
HTTPS/1.1 200 OK
x-amz-id-2: YgIPIfBiKa2bj0KMg95r/0zo3emzU4dzsD4rcKCHQUAdQkf3ShJT00pXUueF6QKo
x-amz-request-id: 236A8905248E5A01
Date: Wed, 01 Mar 2006 12:00:00 GMT
```

```
Location: /colorpictures
Content-Length: 0
Connection: close
Server: AmazonS3
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutBucketWebsite

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Sets the configuration of the website that is specified in the `website` subresource. To configure a bucket as a website, you can add this subresource on the bucket with website configuration information such as the file name of the index document and any redirect rules. For more information, see [Hosting Websites on Amazon S3](#).

This PUT action requires the `S3:PutBucketWebsite` permission. By default, only the bucket owner can configure the website attached to a bucket; however, bucket owners can allow other users to set the website configuration by writing a bucket policy that grants them the `S3:PutBucketWebsite` permission.

To redirect all website requests sent to the bucket's website endpoint, you add a website configuration with the following elements. Because all requests are sent to another website, you don't need to provide index document name for the bucket.

- `WebsiteConfiguration`
- `RedirectAllRequestsTo`
- `HostName`
- `Protocol`

If you want granular control over redirects, you can use the following elements to add routing rules that describe conditions for redirecting requests and information about the redirect destination. In this case, the website configuration must provide an index document for the bucket, because some requests might not be redirected.

- `WebsiteConfiguration`
- `IndexDocument`
- `Suffix`
- `ErrorDocument`

- Key
- RoutingRules
- RoutingRule
- Condition
- HttpStatusCodeReturnedEquals
- KeyPrefixEquals
- Redirect
- Protocol
- HostName
- ReplaceKeyPrefixWith
- ReplaceKeyWith
- HttpRedirectCode

Amazon S3 has a limitation of 50 routing rules per website configuration. If you require more than 50 routing rules, you can use object redirect. For more information, see [Configuring an Object Redirect](#) in the *Amazon S3 User Guide*.

The maximum request length is limited to 128 KB.

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /?website HTTP/1.1
Host: Bucket.s3.amazonaws.com
Content-MD5: ContentMD5
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<WebsiteConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <ErrorDocument>
```

```

    <Key>string</Key>
  </ErrorDocument>
  <IndexDocument>
    <Suffix>string</Suffix>
  </IndexDocument>
  <RedirectAllRequestsTo>
    <HostName>string</HostName>
    <Protocol>string</Protocol>
  </RedirectAllRequestsTo>
  <RoutingRules>
    <RoutingRule>
      <Condition>
        <HttpErrorCodeReturnedEquals>string</HttpErrorCodeReturnedEquals>
        <KeyPrefixEquals>string</KeyPrefixEquals>
      </Condition>
      <Redirect>
        <HostName>string</HostName>
        <HttpRedirectCode>string</HttpRedirectCode>
        <Protocol>string</Protocol>
        <ReplaceKeyPrefixWith>string</ReplaceKeyPrefixWith>
        <ReplaceKeyWith>string</ReplaceKeyWith>
      </Redirect>
    </RoutingRule>
  </RoutingRules>
</WebsiteConfiguration>

```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name.

Required: Yes

Content-MD5

The Base64 encoded 128-bit MD5 digest of the data. You must use this header as a message integrity check to verify that the request body was not corrupted in transit. For more information, see [RFC 1864](#).

For requests made using the AWS Command Line Interface (CLI) or AWS SDKs, this field is calculated automatically.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-sdk-checksum-algorithm

Indicates the algorithm used to create the checksum for the request when you use the SDK. This header will not provide any additional functionality if you don't use the SDK. When you send this header, there must be a corresponding x-amz-checksum or x-amz-trailer header sent. Otherwise, Amazon S3 fails the request with the HTTP status code 400 Bad Request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

If you provide an individual checksum, Amazon S3 ignores any provided ChecksumAlgorithm parameter.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 | XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in XML format.

WebsiteConfiguration

Root level tag for the WebsiteConfiguration parameters.

Required: Yes

ErrorDocument

The name of the error document for the website.

Type: [ErrorDocument](#) data type

Required: No

IndexDocument

The name of the index document for the website.

Type: [IndexDocument](#) data type

Required: No

[RedirectAllRequestsTo](#)

The redirect behavior for every request to this bucket's website endpoint.

Important

If you specify this property, you can't specify any other property.

Type: [RedirectAllRequestsTo](#) data type

Required: No

[RoutingRules](#)

Rules that define when a redirect is applied and the redirect behavior.

Type: Array of [RoutingRule](#) data types

Required: No

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Example 1: Configure bucket as a website (add website configuration)

The following request configures a bucket example.com as a website. The configuration in the request specifies index.html as the index document. It also specifies the optional error document, SomeErrorDocument.html.

```
PUT ?website HTTP/1.1
```



```
Host: example.com.s3.<Region>.amazonaws.com
Content-Length: 256
Date: Thu, 27 Jan 2011 12:00:00 GMT
Authorization: signatureValue

<WebsiteConfiguration xmlns='http://s3.amazonaws.com/doc/2006-03-01/'>
  <IndexDocument>
    <Suffix>index.html</Suffix>
  </IndexDocument>
  <ErrorDocument>
    <Key>SomeErrorDocument.html</Key>
  </ErrorDocument>
</WebsiteConfiguration>
```

Sample Response

This example illustrates one usage of PutBucketWebsite.

```
HTTP/1.1 200 OK
x-amz-id-2: YgIPIfBiKa2bj0KMgUAdQkf3ShJT00pXUueF6QKo
x-amz-request-id: 80CD4368BD211111
Date: Thu, 27 Jan 2011 00:00:00 GMT
Content-Length: 0
Server: AmazonS3
```

Example 2: Configure bucket as a website but redirect all requests

The following request configures a bucket `www.example.com` as a website. However, the configuration specifies that all GET requests for the `www.example.com` bucket's website endpoint will be redirected to host `example.com`. This redirect can be useful when you want to serve requests for both `http://www.example.com` and `http://example.com`, but you want to maintain the website content in only one bucket, in this case, `example.com`.

```
PUT ?website HTTP/1.1
Host: www.example.com.s3.<Region>.amazonaws.com
Content-Length: length-value
Date: Thu, 27 Jan 2011 12:00:00 GMT
Authorization: signatureValue
```

```
<WebsiteConfiguration xmlns='http://s3.amazonaws.com/doc/2006-03-01/'>
  <RedirectAllRequestsTo>
    <HostName>example.com</HostName>
  </RedirectAllRequestsTo>
</WebsiteConfiguration>
```

Example 3: Configure bucket as a website and specify optional redirection rules

Example 1 is the simplest website configuration. It configures a bucket as a website by providing only an index document and an error document. You can further customize the website configuration by adding routing rules that redirect requests for one or more objects. For example, suppose that your bucket contained the following objects:

- index.html
- docs/article1.html
- docs/article2.html

If you decided to rename the folder from docs/ to documents/, you would need to redirect requests for prefix /docs to documents/. For example, a request for docs/article1.html will need to be redirected to documents/article1.html.

In this case, you update the website configuration and add a routing rule as shown in the following request.

```
PUT ?website HTTP/1.1
Host: www.example.com.s3.<Region>.amazonaws.com
Content-Length: length-value
Date: Thu, 27 Jan 2011 12:00:00 GMT
Authorization: signatureValue

<WebsiteConfiguration xmlns='http://s3.amazonaws.com/doc/2006-03-01/'>
  <IndexDocument>
    <Suffix>index.html</Suffix>
  </IndexDocument>
  <ErrorDocument>
    <Key>Error.html</Key>
  </ErrorDocument>
```

```

<RoutingRules>
  <RoutingRule>
    <Condition>
      <KeyPrefixEquals>docs/</KeyPrefixEquals>
    </Condition>
    <Redirect>
      <ReplaceKeyPrefixWith>documents/</ReplaceKeyPrefixWith>
    </Redirect>
  </RoutingRule>
</RoutingRules>
</WebsiteConfiguration>

```

Example 4: Configure a bucket as a website and redirect errors

You can use a routing rule to specify a condition that checks for a specific HTTP error code. When a page request results in this error, you can optionally reroute requests. For example, you might route requests to another host and optionally process the error. The routing rule in the following requests redirects requests to an EC2 instance in the event of an HTTP error 404. For illustration, the redirect also inserts an object key prefix `report-404/` in the redirect. For example, if you request a page `ExamplePage.html` and it results in an HTTP 404 error, the request is routed to a page `report-404/testPage.html` on the specified EC2 instance. If there is no routing rule and the HTTP error 404 occurred, then `Error.html` would be returned.

```

PUT ?website HTTP/1.1
Host: www.example.com.s3.<Region>.amazonaws.com
Content-Length: 580
Date: Thu, 27 Jan 2011 12:00:00 GMT
Authorization: signatureValue

<WebsiteConfiguration xmlns='http://s3.amazonaws.com/doc/2006-03-01/'>
  <IndexDocument>
    <Suffix>index.html</Suffix>
  </IndexDocument>
  <ErrorDocument>
    <Key>Error.html</Key>
  </ErrorDocument>

  <RoutingRules>
    <RoutingRule>
      <Condition>

```

```

    <HttpErrorCodeReturnedEquals>404</HttpErrorCodeReturnedEquals >
  </Condition>
  <Redirect>
    <HostName>ec2-11-22-333-44.compute-1.amazonaws.com</HostName>
    <ReplaceKeyPrefixWith>report-404/</ReplaceKeyPrefixWith>
  </Redirect>
</RoutingRule>
</RoutingRules>
</WebsiteConfiguration>

```

Example 5: Configure a bucket as a website and redirect folder requests to a page

Suppose you have the following pages in your bucket:

- images/photo1.jpg
- images/photo2.jpg
- images/photo3.jpg

Now you want to route requests for all pages with the images/ prefix to go to a single page, errorpage.html. You can add a website configuration to your bucket with the routing rule shown in the following request.

```

PUT ?website HTTP/1.1
Host: www.example.com.s3.<Region>.amazonaws.com
Content-Length: 481
Date: Thu, 27 Jan 2011 12:00:00 GMT
Authorization: signatureValue

<WebsiteConfiguration xmlns='http://s3.amazonaws.com/doc/2006-03-01/'>
  <IndexDocument>
    <Suffix>index.html</Suffix>
  </IndexDocument>
  <ErrorDocument>
    <Key>Error.html</Key>
  </ErrorDocument>

  <RoutingRules>
    <RoutingRule>
      <Condition>

```

```
<KeyPrefixEquals>images/</KeyPrefixEquals>
</Condition>
<Redirect>
  <ReplaceKeyWith>errorpage.html</ReplaceKeyWith>
</Redirect>
</RoutingRule>
</RoutingRules>
</WebsiteConfiguration>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutObject

Service: Amazon S3

Important

End of support notice: As of October 1, 2025, Amazon S3 has discontinued support for Email Grantee Access Control Lists (ACLs). If you attempt to use an Email Grantee ACL in a request after October 1, 2025, the request will receive an HTTP 405 (Method Not Allowed) error.

This change affects the following AWS Regions: US East (N. Virginia), US West (N. California), US West (Oregon), Asia Pacific (Singapore), Asia Pacific (Sydney), Asia Pacific (Tokyo), Europe (Ireland), and South America (São Paulo).

Adds an object to a bucket.

Note

- Amazon S3 never adds partial objects; if you receive a success response, Amazon S3 added the entire object to the bucket. You cannot use PutObject to only update a single piece of metadata for an existing object. You must put the entire object with updated metadata if you want to update some values.
- If your bucket uses the bucket owner enforced setting for Object Ownership, ACLs are disabled and no longer affect permissions. All objects written to the bucket by any account will be owned by the bucket owner.
- **Directory buckets** - For directory buckets, you must make requests for this API operation to the Zonal endpoint. These endpoints support virtual-hosted-style requests in the format `https://amzn-s3-demo-bucket.s3express-zone-id.region-code.amazonaws.com/key-name` . Path-style requests are not supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

Amazon S3 is a distributed system. If it receives multiple write requests for the same object simultaneously, it overwrites all but the last object written. However, Amazon S3 provides features that can modify this behavior:

- **S3 Object Lock** - To prevent objects from being deleted or overwritten, you can use [Amazon S3 Object Lock](#) in the *Amazon S3 User Guide*.

 Note

This functionality is not supported for directory buckets.

- **If-None-Match** - Uploads the object only if the object key name does not already exist in the specified bucket. Otherwise, Amazon S3 returns a 412 Precondition Failed error. If a conflicting operation occurs during the upload, S3 returns a 409 ConditionalRequestConflict response. On a 409 failure, retry the upload.

Expects the * character (asterisk).

For more information, see [Add preconditions to S3 operations with conditional requests](#) in the *Amazon S3 User Guide* or [RFC 7232](#).

 Note

This functionality is not supported for S3 on Outposts.

- **S3 Versioning** - When you enable versioning for a bucket, if Amazon S3 receives multiple write requests for the same object simultaneously, it stores all versions of the objects. For each write request that is made to the same object, Amazon S3 automatically generates a unique version ID of that object being stored in Amazon S3. You can retrieve, replace, or delete any version of the object. For more information about versioning, see [Adding Objects to Versioning-Enabled Buckets](#) in the *Amazon S3 User Guide*. For information about returning the versioning state of a bucket, see [GetBucketVersioning](#).

 Note

This functionality is not supported for directory buckets.

Permissions

- **General purpose bucket permissions** - The following permissions are required in your policies when your `PutObject` request includes specific headers.
 - **s3:PutObject** - To successfully complete the `PutObject` request, you must always have the `s3:PutObject` permission on a bucket to add an object to it.
 - **s3:PutObjectAcl** - To successfully change the objects ACL of your `PutObject` request, you must have the `s3:PutObjectAcl`.
 - **s3:PutObjectTagging** - To successfully set the tag-set with your `PutObject` request, you must have the `s3:PutObjectTagging`.
- **Directory bucket permissions** - To grant access to this API operation on a directory bucket, we recommend that you use the [CreateSession](#) API operation for session-based authorization. Specifically, you grant the `s3express:CreateSession` permission to the directory bucket in a bucket policy or an IAM identity-based policy. Then, you make the `CreateSession` API call on the bucket to obtain a session token. With the session token in your request header, you can make API requests to this operation. After the session token expires, you make another `CreateSession` API call to generate a new session token for use. AWS CLI or SDKs create session and refresh the session token automatically to avoid service interruptions when a session expires. For more information about authorization, see [CreateSession](#).

If the object is encrypted with SSE-KMS, you must also have the `kms:GenerateDataKey` and `kms:Decrypt` permissions in IAM identity-based policies and AWS KMS key policies for the AWS KMS key.

Data integrity with Content-MD5

- **General purpose bucket** - To ensure that data is not corrupted traversing the network, use the `Content-MD5` header. When you use this header, Amazon S3 checks the object against the provided MD5 value and, if they do not match, Amazon S3 returns an error. Alternatively, when the object's ETag is its MD5 digest, you can calculate the MD5 while putting the object to Amazon S3 and compare the returned ETag to the calculated MD5 value.
- **Directory bucket** - This functionality is not supported for directory buckets.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `Bucket-name.s3express-zone-id.region-code.amazonaws.com`.

Errors

- You might receive an `InvalidRequest` error for several reasons. Depending on the reason for the error, you might receive one of the following messages:
 - Cannot specify both a write offset value and user-defined object metadata for existing objects.
 - Checksum Type mismatch occurred, expected checksum Type: sha1, actual checksum Type: crc32c.
 - Request body cannot be empty when 'write offset' is specified.

For more information about related Amazon S3 APIs, see the following:

- [CopyObject](#)
- [DeleteObject](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /Key+ HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-acl: ACL
Cache-Control: CacheControl
Content-Disposition: ContentDisposition
Content-Encoding: ContentEncoding
Content-Language: ContentLanguage
Content-Length: ContentLength
Content-MD5: ContentMD5
Content-Type: ContentType
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-checksum-crc32: ChecksumCRC32
x-amz-checksum-crc32c: ChecksumCRC32C
x-amz-checksum-crc64nvme: ChecksumCRC64NVME
x-amz-checksum-sha1: ChecksumSHA1
```

```

x-amz-checksum-sha256: ChecksumSHA256
x-amz-checksum-sha512: ChecksumSHA512
x-amz-checksum-md5: ChecksumMD5
x-amz-checksum-xxhash64: ChecksumXXHASH64
x-amz-checksum-xxhash3: ChecksumXXHASH3
x-amz-checksum-xxhash128: ChecksumXXHASH128
Expires: Expires
If-Match: IfMatch
If-None-Match: IfNoneMatch
x-amz-grant-full-control: GrantFullControl
x-amz-grant-read: GrantRead
x-amz-grant-read-acp: GrantReadACP
x-amz-grant-write-acp: GrantWriteACP
x-amz-write-offset-bytes: WriteOffsetBytes
x-amz-server-side-encryption: ServerSideEncryption
x-amz-storage-class: StorageClass
x-amz-website-redirect-location: WebsiteRedirectLocation
x-amz-server-side-encryption-customer-algorithm: SSECustomerAlgorithm
x-amz-server-side-encryption-customer-key: SSECustomerKey
x-amz-server-side-encryption-customer-key-MD5: SSECustomerKeyMD5
x-amz-server-side-encryption-aws-kms-key-id: SSEKMSKeyId
x-amz-server-side-encryption-context: SSEKMSEncryptionContext
x-amz-server-side-encryption-bucket-key-enabled: BucketKeyEnabled
x-amz-request-payer: RequestPayer
x-amz-tagging: Tagging
x-amz-object-lock-mode: ObjectLockMode
x-amz-object-lock-retain-until-date: ObjectLockRetainUntilDate
x-amz-object-lock-legal-hold: ObjectLockLegalHoldStatus
x-amz-object-lock-event-hold: ObjectLockEventHold
x-amz-object-lock-event-hold-duration-days: ObjectLockEventHoldDurationDays
x-amz-object-lock-event-hold-duration-years: ObjectLockEventHoldDurationYears
x-amz-expected-bucket-owner: ExpectedBucketOwner

```

Body

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name to which the PUT action was initiated.

Directory buckets - When you use this operation with a directory bucket, you must use virtual-hosted-style requests in the format `Bucket-name.s3express-zone-id.region-code.amazonaws.com`. Path-style requests are not supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must follow the format `bucket-base-name--zone-id--x-s3` (for example, `amzn-s3-demo-bucket--usw2-az1--x-s3`). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form `AccessPointName-AccountId.s3-accesspoint.Region.amazonaws.com`. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

 Note

Object Lambda access points are not supported by directory buckets.

S3 on Outposts - When you use this action with S3 on Outposts, you must direct requests to the S3 on Outposts hostname. The S3 on Outposts hostname takes the form `AccessPointName-AccountId.outpostID.s3-outposts.Region.amazonaws.com`. When you use this action with S3 on Outposts, the destination bucket must be the Outposts access point ARN or the access point alias. For more information about S3 on Outposts, see [What is S3 on Outposts?](#) in the *Amazon S3 User Guide*.

Required: Yes

Cache-Control

Can be used to specify caching behavior along the request/reply chain. For more information, see <http://www.w3.org/Protocols/rfc2616/rfc2616-sec14.html#sec14.9>.

Content-Disposition

Specifies presentational information for the object. For more information, see <https://www.rfc-editor.org/rfc/rfc6266#section-4>.

Content-Encoding

Specifies what content encodings have been applied to the object and thus what decoding mechanisms must be applied to obtain the media-type referenced by the Content-Type header field. For more information, see <https://www.rfc-editor.org/rfc/rfc9110.html#field.content-encoding>.

Content-Language

The language the content is in.

Content-Length

Size of the body in bytes. This parameter is useful when the size of the body cannot be determined automatically. For more information, see <https://www.rfc-editor.org/rfc/rfc9110.html#name-content-length>.

Content-MD5

The Base64 encoded 128-bit MD5 digest of the message (without the headers) according to RFC 1864. This header can be used as a message integrity check to verify that the data is the same data that was originally sent. Although it is optional, we recommend using the Content-MD5 mechanism as an end-to-end integrity check. For more information about REST request authentication, see [REST Authentication](#).

Note

The Content-MD5 or x-amz-sdk-checksum-algorithm header is required for any request to upload an object with a retention period configured using Amazon S3 Object Lock. For more information, see [Uploading objects to an Object Lock enabled bucket](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Content-Type

A standard MIME type describing the format of the contents. For more information, see <https://www.rfc-editor.org/rfc/rfc9110.html#name-content-type>.

Expires

The date and time at which the object is no longer cacheable. For more information, see <https://www.rfc-editor.org/rfc/rfc7234#section-5.3>.

If-Match

Uploads the object only if the ETag (entity tag) value provided during the WRITE operation matches the ETag of the object in S3. If the ETag values do not match, the operation returns a 412 Precondition Failed error.

If a conflicting operation occurs during the upload S3 returns a 409 ConditionalRequestConflict response. On a 409 failure you should fetch the object's ETag and retry the upload.

Expects the ETag value as a string.

For more information about conditional requests, see [RFC 7232](#), or [Conditional requests](#) in the *Amazon S3 User Guide*.

If-None-Match

Uploads the object only if the object key name does not already exist in the bucket specified. Otherwise, Amazon S3 returns a 412 Precondition Failed error.

If a conflicting operation occurs during the upload S3 returns a 409 ConditionalRequestConflict response. On a 409 failure you should retry the upload.

Expects the '*' (asterisk) character.

For more information about conditional requests, see [RFC 7232](#), or [Conditional requests](#) in the *Amazon S3 User Guide*.

Key

Object key for which the PUT action was initiated.

Length Constraints: Minimum length of 1.

Required: Yes

x-amz-acl

The canned ACL to apply to the object. For more information, see [Canned ACL](#) in the *Amazon S3 User Guide*.

When adding a new object, you can use headers to grant ACL-based permissions to individual AWS accounts or to predefined groups defined by Amazon S3. These permissions are then added to the ACL on the object. By default, all objects are private. Only the owner has full access control. For more information, see [Access Control List \(ACL\) Overview](#) and [Managing ACLs Using the REST API](#) in the *Amazon S3 User Guide*.

If the bucket that you're uploading objects to uses the bucket owner enforced setting for S3 Object Ownership, ACLs are disabled and no longer affect permissions. Buckets that use this setting only accept PUT requests that don't specify an ACL or PUT requests that specify bucket owner full control ACLs, such as the `bucket-owner-full-control` canned ACL or an equivalent form of this ACL expressed in the XML format. PUT requests that contain other ACLs (for example, custom grants to certain AWS accounts) fail and return a 400 error with the error code `AccessControlListNotSupported`. For more information, see [Controlling ownership of objects and disabling ACLs](#) in the *Amazon S3 User Guide*.

Note

- This functionality is not supported for directory buckets.
- This functionality is not supported for Amazon S3 on Outposts.

Valid Values: `private` | `public-read` | `public-read-write` | `authenticated-read` | `aws-exec-read` | `bucket-owner-read` | `bucket-owner-full-control`

x-amz-checksum-crc32

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 32-bit CRC32 checksum of the object. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-crc32c

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 32-bit CRC32C checksum of the object. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-crc64nvme

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 64-bit CRC64NVME checksum of the object. The CRC64NVME checksum is always a full object checksum. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-checksum-md5

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 128-bit MD5 digest of the object. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-checksum-sha1

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 160-bit SHA1 digest of the object. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-sha256

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 256-bit SHA256 digest of the object. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-sha512

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 512-bit SHA512 digest of the object. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-checksum-xxhash128

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 128-bit XXHASH128 checksum of the object. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-checksum-xxhash3

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 64-bit XXHASH3 checksum of the object. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-checksum-xxhash64

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 64-bit XXHASH64 checksum of the object. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-grant-full-control

Gives the grantee READ, READ_ACP, and WRITE_ACP permissions on the object.

Note

- This functionality is not supported for directory buckets.
- This functionality is not supported for Amazon S3 on Outposts.

x-amz-grant-read

Allows grantee to read the object data and its metadata.

Note

- This functionality is not supported for directory buckets.
- This functionality is not supported for Amazon S3 on Outposts.

x-amz-grant-read-acp

Allows grantee to read the object ACL.

Note

- This functionality is not supported for directory buckets.
- This functionality is not supported for Amazon S3 on Outposts.

x-amz-grant-write-acp

Allows grantee to write the ACL for the applicable object.

Note

- This functionality is not supported for directory buckets.
- This functionality is not supported for Amazon S3 on Outposts.

x-amz-object-lock-event-hold

Specifies the event hold status to apply to this object. Set to ON to enable or OFF to disable.

Note

This functionality is not supported for directory buckets.

Valid Values: ON | OFF

x-amz-object-lock-event-hold-duration-days

Specifies the event hold duration in days to apply to this object. You cannot specify a duration in both days and years.

Note

This functionality is not supported for directory buckets.

x-amz-object-lock-event-hold-duration-years

Specifies the event hold duration in years to apply to this object. You cannot specify a duration in both days and years.

Note

This functionality is not supported for directory buckets.

x-amz-object-lock-legal-hold

Specifies whether a legal hold will be applied to this object. For more information about S3 Object Lock, see [Object Lock](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: ON | OFF

x-amz-object-lock-mode

The Object Lock mode that you want to apply to this object.

Note

This functionality is not supported for directory buckets.

Valid Values: GOVERNANCE | COMPLIANCE

x-amz-object-lock-retain-until-date

The date and time when you want this object's Object Lock to expire. Must be formatted as a timestamp parameter.

Note

This functionality is not supported for directory buckets.

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

x-amz-sdk-checksum-algorithm

Indicates the algorithm used to create the checksum for the object when you use the SDK. This header will not provide any additional functionality if you don't use the SDK. When you send this header, there must be a corresponding `x-amz-checksum-algorithm` or `x-amz-trailer` header sent. Otherwise, Amazon S3 fails the request with the HTTP status code 400 Bad Request.

For the `x-amz-checksum-algorithm` header, replace *algorithm* with the supported algorithm from the following list:

- CRC32
- CRC32C
- CRC64NVME
- MD5
- SHA1
- SHA256
- SHA512
- XXHASH3
- XXHASH64
- XXHASH128

For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

If the individual checksum value you provide through `x-amz-checksum-algorithm` doesn't match the checksum algorithm you set through `x-amz-sdk-checksum-algorithm`, Amazon S3 fails the request with a BadDigest error.

Note

The Content-MD5 or `x-amz-sdk-checksum-algorithm` header is required for any request to upload an object with a retention period configured using Amazon S3 Object

Lock. For more information, see [Uploading objects to an Object Lock enabled bucket](#) in the *Amazon S3 User Guide*.

For directory buckets, when you use AWS SDKs, CRC32 is the default checksum algorithm that's used for performance.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 | XXHASH64 | XXHASH3 | XXHASH128

x-amz-server-side-encryption

The server-side encryption algorithm that was used when you store this object in Amazon S3 or Amazon FSx.

- **General purpose buckets** - You have four mutually exclusive options to protect data using server-side encryption in Amazon S3, depending on how you choose to manage the encryption keys. Specifically, the encryption key options are Amazon S3 managed keys (SSE-S3), AWS KMS keys (SSE-KMS or DSSE-KMS), and customer-provided keys (SSE-C). Amazon S3 encrypts data with server-side encryption by using Amazon S3 managed keys (SSE-S3) by default. You can optionally tell Amazon S3 to encrypt data at rest by using server-side encryption with other key options. For more information, see [Using Server-Side Encryption](#) in the *Amazon S3 User Guide*.
- **Directory buckets** - For directory buckets, there are only two supported options for server-side encryption: server-side encryption with Amazon S3 managed keys (SSE-S3) (AES256) and server-side encryption with AWS KMS keys (SSE-KMS) (aws:kms). We recommend that the bucket's default encryption uses the desired encryption configuration and you don't override the bucket default encryption in your CreateSession requests or PUT object requests. Then, new objects are automatically encrypted with the desired encryption settings. For more information, see [Protecting data with server-side encryption](#) in the *Amazon S3 User Guide*. For more information about the encryption overriding behaviors in directory buckets, see [Specifying server-side encryption with AWS KMS for new object uploads](#).

In the Zonal endpoint API calls (except [CopyObject](#) and [UploadPartCopy](#)) using the REST API, the encryption request headers must match the encryption settings that are specified in the CreateSession request. You can't override the values of the encryption settings (x-amz-server-side-encryption, x-amz-server-side-encryption-aws-kms-key-id, x-amz-server-side-encryption-context, and x-amz-server-side-encryption-bucket-key-enabled) that are specified in the CreateSession request. You don't need

to explicitly specify these encryption settings values in Zonal endpoint API calls, and Amazon S3 will use the encryption settings values from the `CreateSession` request to protect new objects in the directory bucket.

 Note

When you use the CLI or the AWS SDKs, for `CreateSession`, the session token refreshes automatically to avoid service interruptions when a session expires. The CLI or the AWS SDKs use the bucket's default encryption configuration for the `CreateSession` request. It's not supported to override the encryption settings values in the `CreateSession` request. So in the Zonal endpoint API calls (except [CopyObject](#) and [UploadPartCopy](#)), the encryption request headers must match the default encryption configuration of the directory bucket.

- **S3 access points for Amazon FSx** - When accessing data stored in Amazon FSx file systems using S3 access points, the only valid server side encryption option is `aws:fsx`. All Amazon FSx file systems have encryption configured by default and are encrypted at rest. Data is automatically encrypted before being written to the file system, and automatically decrypted as it is read. These processes are handled transparently by Amazon FSx.

Valid Values: `AES256` | `aws:fsx` | `aws:backup` | `aws:kms` | `aws:kms:dsse`

[x-amz-server-side-encryption-aws-kms-key-id](#)

Specifies the AWS KMS key ID (Key ID, Key ARN, or Key Alias) to use for object encryption. If the KMS key doesn't exist in the same account that's issuing the command, you must use the full Key ARN not the Key ID.

General purpose buckets - If you specify `x-amz-server-side-encryption` with `aws:kms` or `aws:kms:dsse`, this header specifies the ID (Key ID, Key ARN, or Key Alias) of the AWS KMS key to use. If you specify `x-amz-server-side-encryption:aws:kms` or `x-amz-server-side-encryption:aws:kms:dsse`, but do not provide `x-amz-server-side-encryption-aws-kms-key-id`, Amazon S3 uses the AWS managed key (`aws/s3`) to protect the data.

Directory buckets - To encrypt data using SSE-KMS, it's recommended to specify the `x-amz-server-side-encryption` header to `aws:kms`. Then, the `x-amz-server-side-encryption-aws-kms-key-id` header implicitly uses the bucket's default KMS customer managed key ID. If you want to explicitly set the `x-amz-server-side-encryption-aws-kms-key-id` header, it must match the bucket's default customer managed key (using key

ID or ARN, not alias). Your SSE-KMS configuration can only support 1 [customer managed key](#) per directory bucket's lifetime. The [AWS managed key](#) (aws/s3) isn't supported. Incorrect key specification results in an HTTP 400 Bad Request error.

[x-amz-server-side-encryption-bucket-key-enabled](#)

Specifies whether Amazon S3 should use an S3 Bucket Key for object encryption with server-side encryption using AWS Key Management Service (AWS KMS) keys (SSE-KMS).

General purpose buckets - Setting this header to `true` causes Amazon S3 to use an S3 Bucket Key for object encryption with SSE-KMS. Also, specifying this header with a PUT action doesn't affect bucket-level settings for S3 Bucket Key.

Directory buckets - S3 Bucket Keys are always enabled for GET and PUT operations in a directory bucket and can't be disabled. S3 Bucket Keys aren't supported, when you copy SSE-KMS encrypted objects from general purpose buckets to directory buckets, from directory buckets to general purpose buckets, or between directory buckets, through [CopyObject](#), [UploadPartCopy](#), [the Copy operation in Batch Operations](#), or [the import jobs](#). In this case, Amazon S3 makes a call to AWS KMS every time a copy request is made for a KMS-encrypted object.

[x-amz-server-side-encryption-context](#)

Specifies the AWS KMS Encryption Context as an additional encryption context to use for object encryption. The value of this header is a Base64 encoded string of a UTF-8 encoded JSON, which contains the encryption context as key-value pairs. This value is stored as object metadata and automatically gets passed on to AWS KMS for future `GetObject` operations on this object.

General purpose buckets - This value must be explicitly added during `CopyObject` operations if you want an additional encryption context for your object. For more information, see [Encryption context](#) in the *Amazon S3 User Guide*.

Directory buckets - You can optionally provide an explicit encryption context value. The value must match the default encryption context - the bucket Amazon Resource Name (ARN). An additional encryption context value is not supported.

[x-amz-server-side-encryption-customer-algorithm](#)

Specifies the algorithm to use when encrypting the object (for example, AES256).

Note

This functionality is not supported for directory buckets.

[x-amz-server-side-encryption-customer-key](#)

Specifies the customer-provided encryption key for Amazon S3 to use in encrypting data. This value is used to store the object and then it is discarded; Amazon S3 does not store the encryption key. The key must be appropriate for use with the algorithm specified in the `x-amz-server-side-encryption-customer-algorithm` header.

Note

This functionality is not supported for directory buckets.

[x-amz-server-side-encryption-customer-key-MD5](#)

Specifies the 128-bit MD5 digest of the encryption key according to RFC 1321. Amazon S3 uses this header for a message integrity check to ensure that the encryption key was transmitted without error.

Note

This functionality is not supported for directory buckets.

[x-amz-storage-class](#)

By default, Amazon S3 uses the STANDARD Storage Class to store newly created objects. The STANDARD storage class provides high durability and high availability. Depending on performance needs, you can specify a different Storage Class. For more information, see [Storage Classes](#) in the *Amazon S3 User Guide*.

Note

- Directory buckets only support EXPRESS_ONEZONE (the S3 Express One Zone storage class) in Availability Zones and ONEZONE_IA (the S3 One Zone-Infrequent Access storage class) in Dedicated Local Zones.

- Amazon S3 on Outposts only uses the OUTPOSTS Storage Class.

Valid Values: STANDARD | REDUCED_REDUNDANCY | STANDARD_IA | ONEZONE_IA | INTELLIGENT_TIERING | GLACIER | DEEP_ARCHIVE | OUTPOSTS | GLACIER_IR | SNOW | EXPRESS_ONEZONE | FSX_OPENZFS | FSX_ONTAP | AWS_BACKUP_WARM | AWS_BACKUP_LOW_COST_WARM

x-amz-tagging

The tag-set for the object. The tag-set must be encoded as URL Query parameters. (For example, "Key1=Value1")

Note

This functionality is not supported for directory buckets.

x-amz-website-redirect-location

If the bucket is configured as a website, redirects requests for this object to another object in the same bucket or to an external URL. Amazon S3 stores the value of this header in the object metadata. For information about object metadata, see [Object Key and Metadata](#) in the *Amazon S3 User Guide*.

In the following example, the request header sets the redirect to an object (anotherPage.html) in the same bucket:

```
x-amz-website-redirect-location: /anotherPage.html
```

In the following example, the request header sets the object redirect to another website:

```
x-amz-website-redirect-location: http://www.example.com/
```

For more information about website hosting in Amazon S3, see [Hosting Websites on Amazon S3](#) and [How to Configure Website Page Redirects](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

x-amz-write-offset-bytes

Specifies the offset for appending data to existing objects in bytes. The offset must be equal to the size of the existing object being appended to. If no object exists, setting this header to 0 will create a new object.

Note

This functionality is only supported for objects in the Amazon S3 Express One Zone storage class in directory buckets.

Request Body

The request accepts the following binary data.

Body

Response Syntax

```
HTTP/1.1 200
x-amz-expiration: Expiration
ETag: ETag
x-amz-checksum-crc32: ChecksumCRC32
x-amz-checksum-crc32c: ChecksumCRC32C
x-amz-checksum-crc64nvme: ChecksumCRC64NVME
x-amz-checksum-sha1: ChecksumSHA1
x-amz-checksum-sha256: ChecksumSHA256
x-amz-checksum-sha512: ChecksumSHA512
x-amz-checksum-md5: ChecksumMD5
x-amz-checksum-xxhash64: ChecksumXXHASH64
x-amz-checksum-xxhash3: ChecksumXXHASH3
x-amz-checksum-xxhash128: ChecksumXXHASH128
x-amz-checksum-type: ChecksumType
x-amz-server-side-encryption: ServerSideEncryption
x-amz-version-id: VersionId
x-amz-server-side-encryption-customer-algorithm: SSECustomerAlgorithm
x-amz-server-side-encryption-customer-key-MD5: SSECustomerKeyMD5
x-amz-server-side-encryption-aws-kms-key-id: SSEKMSKeyId
x-amz-server-side-encryption-context: SSEKMSEncryptionContext
x-amz-server-side-encryption-bucket-key-enabled: BucketKeyEnabled
```

```
x-amz-object-size: Size  
x-amz-request-charged: RequestCharged
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

ETag

Entity tag for the uploaded object.

General purpose buckets - To ensure that data is not corrupted traversing the network, for objects where the ETag is the MD5 digest of the object, you can calculate the MD5 while putting an object to Amazon S3 and compare the returned ETag to the calculated MD5 value.

Directory buckets - The ETag for the object in a directory bucket isn't the MD5 digest of the object.

x-amz-checksum-crc32

The Base64 encoded, 32-bit CRC32 checksum of the object. This checksum is only present if the checksum was uploaded with the object. When you use an API operation on an object that was uploaded using multipart uploads, this value may not be a direct checksum value of the full object. Instead, it's a calculation based on the checksum values of each individual part. For more information about how checksums are calculated with multipart uploads, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-crc32c

The Base64 encoded, 32-bit CRC32C checksum of the object. This checksum is only present if the checksum was uploaded with the object. When you use an API operation on an object that was uploaded using multipart uploads, this value may not be a direct checksum value of the full object. Instead, it's a calculation based on the checksum values of each individual part. For more information about how checksums are calculated with multipart uploads, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-crc64nvme

The Base64 encoded, 64-bit CRC64NVME checksum of the object. This header is present if the object was uploaded with the CRC64NVME checksum algorithm, or if it was uploaded

without a checksum (and Amazon S3 added the default checksum, CRC64NVME, to the uploaded object). For more information about how checksums are calculated with multipart uploads, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-checksum-md5

The Base64 encoded, 128-bit MD5 digest of the object. This header is present if the object was uploaded with the MD5 checksum algorithm. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-checksum-sha1

The Base64 encoded, 160-bit SHA1 digest of the object. This checksum is only present if the checksum was uploaded with the object. When you use the API operation on an object that was uploaded using multipart uploads, this value may not be a direct checksum value of the full object. Instead, it's a calculation based on the checksum values of each individual part. For more information about how checksums are calculated with multipart uploads, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-sha256

The Base64 encoded, 256-bit SHA256 digest of the object. This checksum is only present if the checksum was uploaded with the object. When you use an API operation on an object that was uploaded using multipart uploads, this value may not be a direct checksum value of the full object. Instead, it's a calculation based on the checksum values of each individual part. For more information about how checksums are calculated with multipart uploads, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-sha512

The Base64 encoded, 512-bit SHA512 digest of the object. This header is present if the object was uploaded with the SHA512 checksum algorithm. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-checksum-type

This header specifies the checksum type of the object, which determines how part-level checksums are combined to create an object-level checksum for multipart objects. For PutObject uploads, the checksum type is always FULL_OBJECT. You can use this header as a data integrity check to verify that the checksum type that is received is the same checksum that was specified. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Valid Values: COMPOSITE | FULL_OBJECT

x-amz-checksum-xxhash128

The Base64 encoded, 128-bit XXHASH128 checksum of the object. This header is present if the object was uploaded with the XXHASH128 checksum algorithm. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-checksum-xxhash3

The Base64 encoded, 64-bit XXHASH3 checksum of the object. This header is present if the object was uploaded with the XXHASH3 checksum algorithm. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-checksum-xxhash64

The Base64 encoded, 64-bit XXHASH64 checksum of the object. This header is present if the object was uploaded with the XXHASH64 checksum algorithm. For more information, see [Checking object integrity in the Amazon S3 User Guide](#).

x-amz-expiration

If the expiration is configured for the object (see [PutBucketLifecycleConfiguration](#)) in the *Amazon S3 User Guide*, the response includes this header. It includes the `expiry-date` and `rule-id` key-value pairs that provide information about object expiration. The value of the `rule-id` is URL-encoded.

Note

Object expiration information is not returned in directory buckets and this header returns the value "NotImplemented" in all responses for directory buckets.

x-amz-object-size

The size of the object in bytes. This value is only be present if you append to an object.

Note

This functionality is only supported for objects in the Amazon S3 Express One Zone storage class in directory buckets.

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

x-amz-server-side-encryption

The server-side encryption algorithm used when you store this object in Amazon S3 or Amazon FSx.

Note

When accessing data stored in Amazon FSx file systems using S3 access points, the only valid server side encryption option is `aws:fsx`.

Valid Values: AES256 | `aws:fsx` | `aws:backup` | `aws:kms` | `aws:kms:dsse`

x-amz-server-side-encryption-aws-kms-key-id

If present, indicates the ID of the KMS key that was used for object encryption.

x-amz-server-side-encryption-bucket-key-enabled

Indicates whether the uploaded object uses an S3 Bucket Key for server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS).

x-amz-server-side-encryption-context

If present, indicates the AWS KMS Encryption Context to use for object encryption. The value of this header is a Base64 encoded string of a UTF-8 encoded JSON, which contains the encryption context as key-value pairs. This value is stored as object metadata and automatically gets passed on to AWS KMS for future `GetObject` operations on this object.

x-amz-server-side-encryption-customer-algorithm

If server-side encryption with a customer-provided encryption key was requested, the response will include this header to confirm the encryption algorithm that's used.

Note

This functionality is not supported for directory buckets.

x-amz-server-side-encryption-customer-key-MD5

If server-side encryption with a customer-provided encryption key was requested, the response will include this header to provide the round-trip message integrity verification of the customer-provided encryption key.

Note

This functionality is not supported for directory buckets.

x-amz-version-id

Version ID of the object.

If you enable versioning for a bucket, Amazon S3 automatically generates a unique version ID for the object being stored. Amazon S3 returns this ID in the response. When you enable versioning for a bucket, if Amazon S3 receives multiple write requests for the same object simultaneously, it stores all of the objects. For more information about versioning, see [Adding Objects to Versioning-Enabled Buckets](#) in the *Amazon S3 User Guide*. For information about returning the versioning state of a bucket, see [GetBucketVersioning](#).

Note

This functionality is not supported for directory buckets.

Errors

EncryptionTypeMismatch

The existing object was created with a different encryption type. Subsequent write requests must include the appropriate encryption parameters in the request or while creating the session.

HTTP Status Code: 400

InvalidRequest

A parameter or header in your request isn't valid. For details, see the description of this API operation.

HTTP Status Code: 400

InvalidWriteOffset

The write offset value that you specified does not match the current object size.

HTTP Status Code: 400

TooManyParts

You have attempted to add more parts than the maximum of 10000 that are allowed for this object. You can use the CopyObject operation to copy this object to another and then add more data to the newly copied object.

HTTP Status Code: 400

Examples

Example 1 for general purpose buckets: Upload an object

The following request stores the `my-image.jpg` file in the `myBucket` bucket.

```
PUT /my-image.jpg HTTP/1.1
Host: myBucket.s3.<Region>.amazonaws.com
Date: Wed, 12 Oct 2009 17:50:00 GMT
Authorization: authorization string
Content-Type: text/plain
Content-Length: 11434
```

```
x-amz-meta-author: Janet
Expect: 100-continue
[11434 bytes of object data]
```

Sample Response for general purpose buckets: Versioning suspended

This example illustrates one usage of PutObject.

```
HTTP/1.1 100 Continue

HTTP/1.1 200 OK
x-amz-id-2: LriYPLdmOdAiIfgSm/F1YsViT1LW94/xUQxMsF7xiEb1a0wiI0Ix1+zbwZ163pt7
x-amz-request-id: 0A49CE4060975EAC
Date: Wed, 12 Oct 2009 17:50:00 GMT
ETag: "1b2cf535f27731c974343645a3985328"
Content-Length: 0
Connection: close
Server: AmazonS3
```

Sample Response for general purpose buckets: Expiration rule created using lifecycle configuration

If an expiration rule that was created on the bucket using lifecycle configuration applies to the object, you get a response with an `x-amz-expiration` header, as shown in the following response. For more information, see [Transitioning Objects: General Considerations](#).

```
HTTP/1.1 100 Continue

HTTP/1.1 200 OK
x-amz-id-2: LriYPLdmOdAiIfgSm/F1YsViT1LW94/xUQxMsF7xiEb1a0wiI0Ix1+zbwZ163pt7
x-amz-request-id: 0A49CE4060975EAC
Date: Wed, 12 Oct 2009 17:50:00 GMT
x-amz-expiration: expiry-date="Fri, 23 Dec 2012 00:00:00 GMT", rule-id="1"
ETag: "1b2cf535f27731c974343645a3985328"
Content-Length: 0
Connection: close
Server: AmazonS3
```


Sample Response for general purpose buckets: Versioning enabled

If the bucket has versioning enabled, the response includes the `x-amz-version-id` header.

```
HTTP/1.1 100 Continue

HTTP/1.1 200 OK
x-amz-id-2: LriYPLdm0dAiIfgSm/F1YsViT1LW94/xUQxMsF7xiEb1a0wiI0Ix1+zbwZ163pt7
x-amz-request-id: 0A49CE4060975EAC
x-amz-version-id: 43jfkodU8493jnFJD9fjj3HHNVfdsQUIFDNsidf038jfdsjGFDSIRp
Date: Wed, 12 Oct 2009 17:50:00 GMT
ETag: "fbacf535f27731c9771645a39863328"
Content-Length: 0
Connection: close
Server: AmazonS3
```

Example 2 for general purpose buckets: Specifying the Reduced Redundancy Storage Class

The following request stores the image, `my-image.jpg`, in the `myBucket` bucket. The request specifies the `x-amz-storage-class` header to request that the object is stored using the `REDUCED_REDUNDANCY` storage class.

```
PUT /my-image.jpg HTTP/1.1
Host: myBucket.s3.<Region>.amazonaws.com
Date: Wed, 12 Oct 2009 17:50:00 GMT
Authorization: authorization string
Content-Type: image/jpeg
Content-Length: 11434
Expect: 100-continue
x-amz-storage-class: REDUCED_REDUNDANCY
```

Sample Response for general purpose buckets

This example illustrates one usage of `PutObject`.

```
HTTP/1.1 100 Continue
```

```
HTTP/1.1 200 OK
x-amz-id-2: LriYPLdm0dAiIfgSm/F1YsViT1LW94/xUQxMsF7xiEb1a0wiI0Ix1+zbwZ163pt7
x-amz-request-id: 0A49CE4060975EAC
Date: Wed, 12 Oct 2009 17:50:00 GMT
ETag: "1b2cf535f27731c974343645a3985328"
Content-Length: 0
Connection: close
Server: AmazonS3
```

Example 3 for general purpose buckets: Uploading an object and specifying access permissions explicitly

The following request stores the `TestObject.txt` file in the `myBucket` bucket. The request specifies various ACL headers to grant permission to AWS accounts that are specified with a canonical user ID and an email address.

```
PUT TestObject.txt HTTP/1.1
Host: myBucket.s3.<Region>.amazonaws.com
x-amz-date: Fri, 13 Apr 2012 05:40:14 GMT
Authorization: authorization string
x-amz-grant-write-acp: id=8a6925ce4adf588a4532142d3f74dd8c71fa124ExampleCanonicalUserID
x-amz-grant-full-control: emailAddress="ExampleUser@amazon.com"
x-amz-grant-write: emailAddress="ExampleUser1@amazon.com",
  emailAddress="ExampleUser2@amazon.com"
Content-Length: 300
Expect: 100-continue
Connection: Keep-Alive

...Object data in the body...
```

Sample Response for general purpose buckets

This example illustrates one usage of `PutObject`.

```
HTTP/1.1 200 OK
x-amz-id-2: RUxG2sZJUfS+ezeAS2i0Xj6w/ST6xqF/8pFNHjTjTrECW56SCAUWGg+7QLVoj1GH
x-amz-request-id: 8D017A90827290BA
Date: Fri, 13 Apr 2012 05:40:25 GMT
```

```
ETag: "dd038b344cf9553547f8b395a814b274"  
Content-Length: 0  
Server: AmazonS3
```

Example 4 for general purpose buckets: Using a canned ACL to set access permissions

The following request stores the `TestObject.txt` file in the `myBucket` bucket. The request uses an `x-amz-acl` header to specify a canned ACL that grants READ permission to the public.

```
PUT TestObject.txt HTTP/1.1  
Host: myBucket.s3.<Region>.amazonaws.com  
x-amz-date: Fri, 13 Apr 2012 05:54:57 GMT  
x-amz-acl: public-read  
Authorization: authorization string  
Content-Length: 300  
Expect: 100-continue  
Connection: Keep-Alive  
  
...Object data in the body...
```

Sample Response for general purpose buckets

This example illustrates one usage of `PutObject`.

```
HTTP/1.1 200 OK  
x-amz-id-2: Yd6PSJxJFQeTYJ/3dD07miqJfVMXXW0S2Hijo3Wfs4bz6oe2QCVXasxXLZdMfASd  
x-amz-request-id: 80DF413BB3D28A25  
Date: Fri, 13 Apr 2012 05:54:59 GMT  
ETag: "dd038b344cf9553547f8b395a814b274"  
Content-Length: 0  
Server: AmazonS3
```

Example 5 for general purpose buckets: Upload an object (Request server-side encryption using a customer-provided encryption key)

This example of an upload object requests server-side encryption and provides an encryption key.

Note

If you have server-side encryption with customer-provided keys (SSE-C) blocked for your general purpose bucket, you will get an HTTP 403 Access Denied error when you specify the SSE-C request headers while writing new data to your bucket. For more information, see [Blocking or unblocking SSE-C for a general purpose bucket](#).

```
PUT /example-object HTTP/1.1
Host: example-bucket.s3.<Region>.amazonaws.com
Accept: */*
Authorization: authorization string
Date: Wed, 28 May 2014 19:31:11 +0000
x-amz-server-side-encryption-customer-key:g0lCfA3Dv40jZz5SQJ1ZukLRFqtI5WorC/8SEEXAMPLE

x-amz-server-side-encryption-customer-key-MD5:ZjQrne1X/iTcskbY2example
x-amz-server-side-encryption-customer-algorithm:AES256
```

Sample Response for general purpose buckets

In the response, Amazon S3 returns the encryption algorithm and MD5 of the encryption key that you specified when uploading the object. The ETag that is returned is not the MD5 of the object.

```
HTTP/1.1 200 OK
x-amz-id-2: 7qoYGN7uMuFuYS6m7a4lshH6in+hccE+4DXPmDZ7C9KqucjinZC1gI5mshai6fbMG
x-amz-request-id: 06437EDD40C407C7
Date: Wed, 28 May 2014 19:31:12 GMT
x-amz-server-side-encryption-customer-algorithm: AES256
x-amz-server-side-encryption-customer-key-MD5: ZjQrne1X/iTcskbY2example
ETag: "ae89237c20e759c5f479ece02c642f59"
```

Example 6 for general purpose buckets: Upload an object and specify tags

This example of an upload object request specifies the optional `x-amz-tagging` header to add tags to the object.

After the object is created, Amazon S3 stores the specified object tags in the tagging subresource that is associated with the object. For more information about tagging, see [Object Tagging and Access Control Policies](#) in the *Amazon S3 User Guide*.

```
PUT /example-object HTTP/1.1
Host: example-bucket.s3.<Region>.amazonaws.com
Accept: */*
Authorization: authorization string
Date: Thu, 22 Sep 2016 21:58:13 GMT
x-amz-tagging: tag1=value1&tag2=value2

[... bytes of object data]
```

Sample Response for general purpose buckets

This example illustrates one usage of PutObject.

```
HTTP/1.1 200 OK
x-amz-id-2: 7qoYGN7uMuFuYS6m7a4lszH6in+hccE+4DXPmDZ7C9KqucjnZC1gI5mshai6fbMG
x-amz-request-id: 06437EDD40C407C7
Date: Thu, 22 Sep 2016 21:58:17 GMT
```

Example 7 for general purpose buckets: Upload an object and specify the checksum algorithm

This example of an upload object request specifies the additional checksum algorithm to use to verify the content of the object. For more information about using additional checksums, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

```
PUT /example-object HTTP/1.1
Host: example-bucket.s3.<Region>.amazonaws.com
x-amz-date: Mon, 22 Mar 2021 23:00:00 GMT
Authorization: authorization string
Content-Length: 268435456
x-amz-checksum-sha256: 0ea4be78f6c3948588172edc6d8789ffe3cec461f385e0ac447e581731c429b5
[268435456 bytes of object data in the body]
```

Sample Response for general purpose buckets

This example illustrates one usage of PutObject.

```
HTTP/1.1 200 OK
x-amz-id-2: 7qoYGN7uMuFuYS6m7a4lszH6in+hccE+4DXPmDZ7C9KqucJnZC1gI5mshai6fbMG
x-amz-request-id: 49CFA2051300FBE9
Date: Mon, 22 Mar 2021 23:00:12 GMT
```

Example 8 for directory buckets: Upload an object and append to it

The following request creates the `my-application.log` file in the `mybucket` bucket, and appends to it afterwards.

```
PUT /my-application.log HTTP/1.1
Host: mybucket--usw2-az1--x-s3
Date: Fri, 22 Nov 2024 17:50:00 GMT
Authorization: authorization string
Content-Type: text/plain
Content-Length: 1048576
[1048576 bytes of object data]

PUT /my-application.log HTTP/1.1
Host: mybucket--usw2-az1--x-s3
Date: Fri, 22 Nov 2024 17:50:00 GMT
Authorization: authorization string
Content-Type: text/plain
Content-Length: 524288
x-amz-write-offset-bytes: 1048576
[524288 bytes of object data]
```

Sample Response for directory buckets

This example illustrates one usage of PutObject.

```
HTTP/1.1 200 OK
```

```
x-amz-request-id: 06437EDD40C407C7
x-amz-id-2: 7qoYGN7uMuFuYS6m7a4lszH6in+hccE+4DXPmDZ7C9KqucjnZC1gI5mshai6fbMG
etag: "ae89237c20e759c5f479ece02c642f59"
x-amz-object-size: 1572864
```

Constraints (AI generated)

This content was generated using AI technology that has in-depth knowledge of AWS API operations. [Provide feedback on this content.](#)

For Encryption parameters:

- You must not pass parameters that don't apply to the type of encryption that you're setting. For example, you cannot pass both `SSECustomerKey`, which is used for server-side encryption with customer-provided keys (SSE-C), and `SSEKMSKeyId`, which is used for server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS), in the same request.
- When using server-side encryption with customer-provided keys (SSE-C), you must pass all three parameters: `SSECustomerAlgorithm`, `SSECustomerKey`, and `SSECustomerKeyMD5`. In addition, the value of `SSECustomerAlgorithm` must be `AES256`.
- When using server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS), you must set the value of `ServerSideEncryption` to `aws:kms`. You can also use `SSEKMSKeyId` to specify the ARN for a specific customer managed KMS encryption key. Additionally, you cannot pass `SSECustomerAlgorithm` when using SSE-KMS.

For Object Lock parameters:

- Object Lock parameters must be paired. If you specify `ObjectLockMode`, you must also specify `ObjectLockRetainUntilDate`, and vice versa. Either specify both parameters or neither one.

For Checksum parameters:

- When you specify a particular `ChecksumAlgorithm` value, you can only pass the checksum parameter that corresponds to that algorithm. For example, if you specify the value `CRC32` for `ChecksumAlgorithm`, you must only pass `ChecksumCRC32` for the CRC-32 algorithm. In this case, you cannot pass any of the following parameters, which are specific to the CRC-32C (`CRC32C`), SHA-1 (`SHA1`), and SHA-256 (`SHA256`) checksum algorithms respectively: `ChecksumCRC32C`, `ChecksumSHA1`, or `ChecksumSHA256`.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutObjectAcl

Service: Amazon S3

Important

End of support notice: As of October 1, 2025, Amazon S3 has discontinued support for Email Grantee Access Control Lists (ACLs). If you attempt to use an Email Grantee ACL in a request after October 1, 2025, the request will receive an HTTP 405 (Method Not Allowed) error.

This change affects the following AWS Regions: US East (N. Virginia), US West (N. California), US West (Oregon), Asia Pacific (Singapore), Asia Pacific (Sydney), Asia Pacific (Tokyo), Europe (Ireland), and South America (São Paulo).

Note

This operation is not supported for directory buckets.

Uses the `acl` subresource to set the access control list (ACL) permissions for a new or existing object in an S3 bucket. You must have the `WRITE_ACP` permission to set the ACL of an object. For more information, see [What permissions can I grant?](#) in the *Amazon S3 User Guide*.

This functionality is not supported for Amazon S3 on Outposts.

Depending on your application needs, you can choose to set the ACL on an object using either the request body or the headers. For example, if you have an existing application that updates a bucket ACL using the request body, you can continue to use that approach. For more information, see [Access Control List \(ACL\) Overview](#) in the *Amazon S3 User Guide*.

Important

If your bucket uses the bucket owner enforced setting for S3 Object Ownership, ACLs are disabled and no longer affect permissions. You must use policies to grant access to your bucket and the objects in it. Requests to set ACLs or update ACLs fail and return the `AccessControlListNotSupported` error code. Requests to read ACLs are still supported. For more information, see [Controlling object ownership](#) in the *Amazon S3 User Guide*.

Permissions

You can set access permissions using one of the following methods:

- Specify a canned ACL with the `x-amz-acl` request header. Amazon S3 supports a set of predefined ACLs, known as canned ACLs. Each canned ACL has a predefined set of grantees and permissions. Specify the canned ACL name as the value of `x-amz-acl`. If you use this header, you cannot use other access control-specific headers in your request. For more information, see [Canned ACL](#).
- Specify access permissions explicitly with the `x-amz-grant-read`, `x-amz-grant-read-acp`, `x-amz-grant-write-acp`, and `x-amz-grant-full-control` headers. When using these headers, you specify explicit access permissions and grantees (AWS accounts or Amazon S3 groups) who will receive the permission. If you use these ACL-specific headers, you cannot use `x-amz-acl` header to set a canned ACL. These parameters map to the set of permissions that Amazon S3 supports in an ACL. For more information, see [Access Control List \(ACL\) Overview](#).

You specify each grantee as a `type=value` pair, where the type is one of the following:

- `id` – if the value specified is the canonical user ID of an AWS account
- `uri` – if you are granting permissions to a predefined group
- `emailAddress` – if the value specified is the email address of an AWS account

Note

Using email addresses to specify a grantee is only supported in the following AWS Regions:

- US East (N. Virginia)
- US West (N. California)
- US West (Oregon)
- Asia Pacific (Singapore)
- Asia Pacific (Sydney)
- Asia Pacific (Tokyo)
- Europe (Ireland)
- South America (São Paulo)

For a list of all the Amazon S3 supported Regions and endpoints, see [Regions and Endpoints](#) in the AWS General Reference.

For example, the following `x-amz-grant-read` header grants list objects permission to the two AWS accounts identified by their email addresses.

```
x-amz-grant-read: emailAddress="xyz@amazon.com",  
emailAddress="abc@amazon.com"
```

You can use either a canned ACL or specify access permissions explicitly. You cannot do both.

Grantee Values

You can specify the person (grantee) to whom you're assigning access rights (using request elements) in the following ways. For examples of how to specify these grantee values in JSON format, see the AWS CLI example in [Enabling Amazon S3 server access logging](#) in the *Amazon S3 User Guide*.

- By the person's ID:

```
<Grantee xmlns:xsi="http://www.w3.org/2001/XMLSchema-  
instance" xsi:type="CanonicalUser"><ID><>ID<></  
ID><DisplayName><>GranteesEmail<></DisplayName> </Grantee>
```

`DisplayName` is optional and ignored in the request.

- By URI:

```
<Grantee xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"  
xsi:type="Group"><URI><>http://acs.amazonaws.com/groups/global/  
AuthenticatedUsers<></URI></Grantee>
```

- By Email address:

```
<Grantee xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"  
xsi:type="AmazonCustomerByEmail"><EmailAddress><>Grantees@email.com<></  
EmailAddress></Grantee>
```

The grantee is resolved to the `CanonicalUser` and, in a response to a GET Object acl request, appears as the `CanonicalUser`.

Note

Using email addresses to specify a grantee is only supported in the following AWS Regions:

- US East (N. Virginia)
- US West (N. California)
- US West (Oregon)
- Asia Pacific (Singapore)
- Asia Pacific (Sydney)
- Asia Pacific (Tokyo)
- Europe (Ireland)
- South America (São Paulo)

For a list of all the Amazon S3 supported Regions and endpoints, see [Regions and Endpoints](#) in the AWS General Reference.

Versioning

The ACL of an object is set at the object version level. By default, PUT sets the ACL of the current version of an object. To set the ACL of a different version, use the `versionId` subresource.

The following operations are related to `PutObjectAcl`:

- [CopyObject](#)
- [GetObject](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /{Key+}?acl&versionId=VersionId HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-acl: ACL
Content-MD5: ContentMD5
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-grant-full-control: GrantFullControl
x-amz-grant-read: GrantRead
x-amz-grant-read-acp: GrantReadACP
x-amz-grant-write: GrantWrite
x-amz-grant-write-acp: GrantWriteACP
x-amz-request-payer: RequestPayer
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<AccessControlPolicy xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <AccessControlList>
    <Grant>
      <Grantee>
        <DisplayName>string</DisplayName>
        <EmailAddress>string</EmailAddress>
        <ID>string</ID>
        <xsi:type>string</xsi:type>
        <URI>string</URI>
      </Grantee>
      <Permission>string</Permission>
    </Grant>
  </AccessControlList>
  <Owner>
    <DisplayName>string</DisplayName>
    <ID>string</ID>
  </Owner>
</AccessControlPolicy>
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name that contains the object to which you want to attach the ACL.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form *AccessPointName-AccountId.s3-accesspoint.Region.amazonaws.com*. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

S3 on Outposts - When you use this action with S3 on Outposts, you must direct requests to the S3 on Outposts hostname. The S3 on Outposts hostname takes the form *AccessPointName-AccountId.outpostID.s3-outposts.Region.amazonaws.com*. When you use this action with S3 on Outposts, the destination bucket must be the Outposts access point ARN or the access point alias. For more information about S3 on Outposts, see [What is S3 on Outposts?](#) in the *Amazon S3 User Guide*.

Required: Yes

Content-MD5

The Base64 encoded 128-bit MD5 digest of the data. This header must be used as a message integrity check to verify that the request body was not corrupted in transit. For more information, go to [RFC 1864](#).

For requests made using the AWS Command Line Interface (CLI) or AWS SDKs, this field is calculated automatically.

Key

Key for which the PUT action was initiated.

Length Constraints: Minimum length of 1.

Required: Yes

versionId

Version ID used to reference a specific version of the object.

Note

This functionality is not supported for directory buckets.

x-amz-acl

The canned ACL to apply to the object. For more information, see [Canned ACL](#).

Valid Values: private | public-read | public-read-write | authenticated-read
| aws-exec-read | bucket-owner-read | bucket-owner-full-control

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-grant-full-control

Allows grantee the read, write, read ACP, and write ACP permissions on the bucket.

This functionality is not supported for Amazon S3 on Outposts.

x-amz-grant-read

Allows grantee to list the objects in the bucket.

This functionality is not supported for Amazon S3 on Outposts.

x-amz-grant-read-acp

Allows grantee to read the bucket ACL.

This functionality is not supported for Amazon S3 on Outposts.

x-amz-grant-write

Allows grantee to create new objects in the bucket.

For the bucket and object owners of existing objects, also allows deletions and overwrites of those objects.

x-amz-grant-write-acp

Allows grantee to write the ACL for the applicable bucket.

This functionality is not supported for Amazon S3 on Outposts.

[x-amz-request-payer](#)

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

[x-amz-sdk-checksum-algorithm](#)

Indicates the algorithm used to create the checksum for the object when you use the SDK. This header will not provide any additional functionality if you don't use the SDK. When you send this header, there must be a corresponding `x-amz-checksum` or `x-amz-trailer` header sent. Otherwise, Amazon S3 fails the request with the HTTP status code 400 Bad Request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

If you provide an individual checksum, Amazon S3 ignores any provided `ChecksumAlgorithm` parameter.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 | XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in XML format.

[AccessControlPolicy](#)

Root level tag for the `AccessControlPolicy` parameters.

Required: Yes

Grants

A list of grants.

Type: Array of [Grant](#) data types

Required: No

Owner

Container for the bucket owner's display name and ID.

Type: [Owner](#) data type

Required: No

Response Syntax

```
HTTP/1.1 200
x-amz-request-charged: RequestCharged
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

Errors

NoSuchKey

The specified key does not exist.

HTTP Status Code: 404

Examples

Sample Request

The following request grants access permission to an existing object. The request specifies the ACL in the body. In addition to granting full control to the object owner, the XML specifies full control to an AWS account identified by its canonical user ID.

```
PUT /my-image.jpg?acl HTTP/1.1
Host: bucket.s3.<Region>.amazonaws.com
Date: Wed, 28 Oct 2009 22:32:00 GMT
Authorization: authorization string
Content-Length: 124

<AccessControlPolicy>
  <Owner>
    <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeebf76c078efc7c6caea54ba06a</ID>
    <DisplayName>CustomersName@amazon.com</DisplayName>
  </Owner>
  <AccessControlList>
    <Grant>
      <Grantee xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:type="CanonicalUser">
        <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeeExampleCanonicalUserID</ID>
        <DisplayName>CustomerName@amazon.com</DisplayName>
      </Grantee>
      <Permission>FULL_CONTROL</Permission>
    </Grant>
  </AccessControlList>
</AccessControlPolicy>
```

Sample Response

The following shows a sample response when versioning on the bucket is enabled.

```
HTTP/1.1 200 OK
x-amz-id-2: eftixk72aD6Ap51T9AS1ed40pIszj7UDNEHGran
x-amz-request-id: 318BC8BC148832E5
x-amz-version-id: 3/L4kqtJlcpXrof3vjVBH40Nr8X8gdRQBpUMLUo
Date: Wed, 28 Oct 2009 22:32:00 GMT
Last-Modified: Sun, 1 Jan 2006 12:00:00 GMT
Content-Length: 0
Connection: close
Server: AmazonS3
```

Sample Request: Setting the ACL of a specified object version

The following request sets the ACL on the specified version of the object.

```
PUT /my-image.jpg?acl&versionId=3HL4kqtJlcpXroDTDmJ+rmSpXd3dIbrHY+MTRCxf3vjVBH40Nrjfkfkd
HTTP/1.1
Host: bucket.s3.<Region>.amazonaws.com
Date: Wed, 28 Oct 2009 22:32:00 GMT
Authorization: authorization string
Content-Length: 124

<AccessControlPolicy>
  <Owner>
    <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeebf76c078efc7c6caea54ba06a</ID>
    <DisplayName>mtd@amazon.com</DisplayName>
  </Owner>
  <AccessControlList>
    <Grant>
      <Grantee xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:type="CanonicalUser">
        <ID>75aa57f09aa0c8caeab4f8c24e99d10f8e7faeebf76c078efc7c6caea54ba06a</ID>
        <DisplayName>mtd@amazon.com</DisplayName>
      </Grantee>
      <Permission>FULL_CONTROL</Permission>
    </Grant>
  </AccessControlList>
```

```
</AccessControlPolicy>
```

Sample Response

This example illustrates one usage of PutObjectAcl.

```
HTTP/1.1 200 OK
x-amz-id-2: eftixk72aD6Ap51u8yU9AS1ed40pIszj7UDNEHGran
x-amz-request-id: 318BC8BC148832E5
x-amz-version-id: 3/L4kqtJlcpXro3vjVBH40Nr8X8gdRQBpUMLUo
Date: Wed, 28 Oct 2009 22:32:00 GMT
Last-Modified: Sun, 1 Jan 2006 12:00:00 GMT
Content-Length: 0
Connection: close
Server: AmazonS3
```

Sample Request: Access permissions specified using headers

The following request sets the ACL on the specified version of the object.

```
PUT ExampleObject.txt?acl HTTP/1.1
Host: examplebucket.s3.<Region>.amazonaws.com
x-amz-acl: public-read
Accept: */*
Authorization: authorization string
Host: s3.amazonaws.com
Connection: Keep-Alive
```

Sample Response

This example illustrates one usage of PutObjectAcl.

```
HTTP/1.1 200 OK
x-amz-id-2: w5YegkbG6ZDsje4WK56RWPxNQHIQ0CjrjyRVFZhEJI9E3kbabXnB09w5G7Dmxsgk
x-amz-request-id: C13B2827BD8455B1
Date: Sun, 29 Apr 2012 23:24:12 GMT
```

```
Content-Length: 0
Server: AmazonS3
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutObjectAnnotation

Service: Amazon S3

Attaches an annotation to an Amazon S3 object. An annotation is a named payload of 1 byte to 1 MiB that you can associate with a specific object or object version. Each object can have up to 1,000 annotations.

For annotation naming rules and restrictions, see [Annotation naming guidelines](#) in the *Amazon S3 User Guide*.

Annotations inherit the encryption of their parent object. For objects without server-side encryption, annotations are encrypted with SSE-S3 (the default for new objects). Objects encrypted with SSE-C cannot have annotations.

To use this operation, you must have the `s3:PutObjectAnnotation` permission. If the bucket has Requester Pays enabled, you must include the `x-amz-request-payer` header.

Note

Annotations are not supported by the following features: S3 Inventory Reports, API Gateway, S3 Storage Lens, Amazon S3 File Gateway, Amazon FSx, S3 on Outposts, and S3 Express One Zone (directory buckets).

The following operations are related to PutObjectAnnotation:

- [GetObjectAnnotation](#)
- [ListObjectAnnotations](#)
- [DeleteObjectAnnotation](#)

Request Syntax

```
PUT /{Key+}?annotation&annotationName=AnnotationName&versionId=VersionId HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-object-if-match: ObjectIfMatch
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-checksum-crc32: ChecksumCRC32
x-amz-checksum-crc32c: ChecksumCRC32C
x-amz-checksum-crc64nvme: ChecksumCRC64NVME
```

```
x-amz-checksum-sha1: ChecksumSHA1
x-amz-checksum-sha256: ChecksumSHA256
x-amz-checksum-sha512: ChecksumSHA512
x-amz-checksum-md5: ChecksumMD5
x-amz-checksum-xxhash64: ChecksumXXHASH64
x-amz-checksum-xxhash3: ChecksumXXHASH3
x-amz-checksum-xxhash128: ChecksumXXHASH128
Content-MD5: ContentMD5
x-amz-request-payer: RequestPayer
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

AnnotationPayload

URI Request Parameters

The request uses the following URI parameters.

annotationName

The name of the annotation.

Length Constraints: Minimum length of 1. Maximum length of 512 bytes.

Required: Yes

Bucket

The name of the bucket that contains the object.

Required: Yes

Content-MD5

Base64-encoded MD5 digest of the message.

Key

The object key.

Length Constraints: Minimum length of 1.

Required: Yes

versionId

The version ID of the object to attach the annotation to.

[x-amz-checksum-crc32](#)

Base64-encoded CRC32 checksum of the annotation payload.

[x-amz-checksum-crc32c](#)

Base64-encoded CRC32C checksum of the annotation payload.

[x-amz-checksum-crc64nvme](#)

Base64-encoded CRC64NVME checksum of the annotation payload.

[x-amz-checksum-md5](#)

Base64-encoded MD5 checksum of the annotation payload.

[x-amz-checksum-sha1](#)

Base64-encoded SHA1 checksum of the annotation payload.

[x-amz-checksum-sha256](#)

Base64-encoded SHA256 checksum of the annotation payload.

[x-amz-checksum-sha512](#)

Base64-encoded SHA512 checksum of the annotation payload.

[x-amz-checksum-xxhash128](#)

Base64-encoded XXHASH128 checksum of the annotation payload.

[x-amz-checksum-xxhash3](#)

Base64-encoded XXHASH3 checksum of the annotation payload.

[x-amz-checksum-xxhash64](#)

Base64-encoded XXHASH64 checksum of the annotation payload.

[x-amz-expected-bucket-owner](#)

The account ID of the expected bucket owner. If the bucket is owned by a different account, the request fails with an HTTP 403 (Access Denied) error.

[x-amz-object-if-match](#)

If specified, the operation only succeeds if the object's ETag matches the provided value.

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

x-amz-sdk-checksum-algorithm

The checksum algorithm to use. Supported values: CRC32, CRC32C, CRC64NVME, SHA1, SHA256, SHA512, MD5, XXHASH64, XXHASH3, XXHASH128.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 | XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following binary data.

AnnotationPayload

Response Syntax

```
HTTP/1.1 200
x-amz-object-version-id: ObjectVersionId
ETag: ETag
x-amz-checksum-crc32: ChecksumCRC32
x-amz-checksum-crc32c: ChecksumCRC32C
x-amz-checksum-crc64nvme: ChecksumCRC64NVME
x-amz-checksum-sha1: ChecksumSHA1
x-amz-checksum-sha256: ChecksumSHA256
x-amz-checksum-sha512: ChecksumSHA512
```

```

x-amz-checksum-md5: ChecksumMD5
x-amz-checksum-xxhash64: ChecksumXXHASH64
x-amz-checksum-xxhash3: ChecksumXXHASH3
x-amz-checksum-xxhash128: ChecksumXXHASH128
x-amz-checksum-type: ChecksumType
x-amz-server-side-encryption: ServerSideEncryption
x-amz-request-charged: RequestCharged
<?xml version="1.0" encoding="UTF-8"?>
<PutObjectAnnotationOutput>
  <Key>string</Key>
  <AnnotationName>string</AnnotationName>
</PutObjectAnnotationOutput>

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

ETag

The entity tag of the annotation.

x-amz-checksum-crc32

The CRC32 checksum of the stored annotation.

x-amz-checksum-crc32c

The CRC32C checksum of the stored annotation.

x-amz-checksum-crc64nvme

The CRC64NVME checksum of the stored annotation.

x-amz-checksum-md5

The MD5 checksum of the stored annotation.

x-amz-checksum-sha1

The SHA1 checksum of the stored annotation.

x-amz-checksum-sha256

The SHA256 checksum of the stored annotation.

x-amz-checksum-sha512

The SHA512 checksum of the stored annotation.

x-amz-checksum-type

The type of checksum used.

Valid Values: COMPOSITE | FULL_OBJECT

x-amz-checksum-xxhash128

The XXHASH128 checksum of the stored annotation.

x-amz-checksum-xxhash3

The XXHASH3 checksum of the stored annotation.

x-amz-checksum-xxhash64

The XXHASH64 checksum of the stored annotation.

x-amz-object-version-id

The version ID of the object that the annotation was attached to.

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

x-amz-server-side-encryption

The server-side encryption algorithm used to encrypt the annotation.

Valid Values: AES256 | aws:fsx | aws:backup | aws:kms | aws:kms:dsse

The following data is returned in XML format by the service.

PutObjectAnnotationOutput

Root level tag for the PutObjectAnnotationOutput parameters.

Required: Yes

AnnotationName

The name of the annotation.

Type: String

Key

The object key.

Type: String

Length Constraints: Minimum length of 1.

Errors

AnnotationLimitExceeded

The request would exceed the maximum number of annotations allowed per object.

HTTP Status Code: 400

AnnotationNameTooLong

The annotation name exceeds 512 bytes.

HTTP Status Code: 400

InvalidAnnotationName

The annotation name you provided is invalid.

HTTP Status Code: 400

InvalidRequest

A parameter or header in your request isn't valid. For details, see the description of this API operation.

HTTP Status Code: 400

NoSuchBucket

The specified bucket does not exist.

HTTP Status Code: 404

NoSuchKey

The specified key does not exist.

HTTP Status Code: 404

UnsupportedMediaType

The annotation payload is not valid UTF-8 encoded text.

HTTP Status Code: 415

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutObjectLegalHold

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Applies a legal hold configuration to the specified object. For more information, see [Locking Objects](#).

This functionality is not supported for Amazon S3 on Outposts.

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /{Key+}?legal-hold&versionId=VersionId HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-request-payer: RequestPayer
Content-MD5: ContentMD5
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<LegalHold xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Status>string</Status>
</LegalHold>
```

URI Request Parameters

The request uses the following URI parameters.

[Bucket](#)

The bucket name containing the object that you want to place a legal hold on.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form *AccessPointName-AccountId.s3-accesspoint.Region.amazonaws.com*. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

Required: Yes

Content-MD5

The MD5 hash for the request body.

For requests made using the AWS Command Line Interface (CLI) or AWS SDKs, this field is calculated automatically.

Key

The key name for the object that you want to place a legal hold on.

Length Constraints: Minimum length of 1.

Required: Yes

versionId

The version ID of the object that you want to place a legal hold on.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

x-amz-sdk-checksum-algorithm

Indicates the algorithm used to create the checksum for the object when you use the SDK. This header will not provide any additional functionality if you don't use the SDK. When you send this header, there must be a corresponding `x-amz-checksum` or `x-amz-trailer` header sent. Otherwise, Amazon S3 fails the request with the HTTP status code 400 Bad Request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

If you provide an individual checksum, Amazon S3 ignores any provided `ChecksumAlgorithm` parameter.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 | XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in XML format.

LegalHold

Root level tag for the LegalHold parameters.

Required: Yes

Status

Indicates whether the specified object has a legal hold in place.

Type: String

Valid Values: ON | OFF

Required: No

Response Syntax

```
HTTP/1.1 200
x-amz-request-charged: RequestCharged
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)

- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutObjectLockConfiguration

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Places an Object Lock configuration on the specified bucket. The rule specified in the Object Lock configuration will be applied by default to every new object placed in the specified bucket. For more information, see [Locking Objects](#).

Note

- The `DefaultRetention` settings require both a mode and a period.
- The `DefaultRetention` period can be either `Days` or `Years` but you must select one. You cannot specify `Days` and `Years` at the same time.
- You can enable Object Lock for new or existing buckets. For more information, see [Configuring Object Lock](#).

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /?object-lock HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-request-payer: RequestPayer
x-amz-bucket-object-lock-token: Token
Content-MD5: ContentMD5
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
```

```
<ObjectLockConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <ObjectLockEnabled>string</ObjectLockEnabled>
  <Rule>
    <DefaultRetention>
      <Days>integer</Days>
      <DefaultEventHold>
        <Days>integer</Days>
        <Years>integer</Years>
      </DefaultEventHold>
      <Mode>string</Mode>
      <Years>integer</Years>
    </DefaultRetention>
  </Rule>
</ObjectLockConfiguration>
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket whose Object Lock configuration you want to create or replace.

Required: Yes

Content-MD5

The MD5 hash for the request body.

For requests made using the AWS Command Line Interface (CLI) or AWS SDKs, this field is calculated automatically.

x-amz-bucket-object-lock-token

A token to allow Object Lock to be enabled for an existing bucket.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3

bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

x-amz-sdk-checksum-algorithm

Indicates the algorithm used to create the checksum for the object when you use the SDK. This header will not provide any additional functionality if you don't use the SDK. When you send this header, there must be a corresponding x-amz-checksum or x-amz-trailer header sent. Otherwise, Amazon S3 fails the request with the HTTP status code 400 Bad Request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

If you provide an individual checksum, Amazon S3 ignores any provided ChecksumAlgorithm parameter.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 | XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in XML format.

ObjectLockConfiguration

Root level tag for the ObjectLockConfiguration parameters.

Required: Yes

ObjectLockEnabled

Indicates whether this bucket has an Object Lock configuration enabled. Enable ObjectLockEnabled when you apply ObjectLockConfiguration to a bucket.

Type: String

Valid Values: Enabled

Required: No

Rule

Specifies the Object Lock rule for the specified object. Enable the this rule when you apply `ObjectLockConfiguration` to a bucket. Bucket settings require both a mode and a period. The period can be either `Days` or `Years` but you must select one. You cannot specify `Days` and `Years` at the same time.

Type: [ObjectLockRule](#) data type

Required: No

Response Syntax

```
HTTP/1.1 200
x-amz-request-charged: RequestCharged
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

[x-amz-request-charged](#)

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: `requester`

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutObjectRetention

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Places an Object Retention configuration on an object. For more information, see [Locking Objects](#). Users or accounts require the `s3:PutObjectRetention` permission in order to place an Object Retention configuration on objects. Bypassing a Governance Retention configuration requires the `s3:BypassGovernanceRetention` permission.

This functionality is not supported for Amazon S3 on Outposts.

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /{Key+}?retention&versionId=VersionId HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-request-payer: RequestPayer
x-amz-bypass-governance-retention: BypassGovernanceRetention
Content-MD5: ContentMD5
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<Retention xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Mode>string</Mode>
  <RetainUntilDate>timestamp</RetainUntilDate>
  <EventHold>string</EventHold>
  <EventHoldDuration>
    <Days>integer</Days>
    <Years>integer</Years>
  </EventHoldDuration>
```


</Retention>

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name that contains the object you want to apply this Object Retention configuration to.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form *AccessPointName-AccountId.s3-accesspoint.Region.amazonaws.com*. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

Required: Yes

Content-MD5

The MD5 hash for the request body.

For requests made using the AWS Command Line Interface (CLI) or AWS SDKs, this field is calculated automatically.

Key

The key name for the object that you want to apply this Object Retention configuration to.

Length Constraints: Minimum length of 1.

Required: Yes

versionId

The version ID for the object that you want to apply this Object Retention configuration to.

x-amz-bypass-governance-retention

Indicates whether this action should bypass Governance-mode restrictions.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

x-amz-sdk-checksum-algorithm

Indicates the algorithm used to create the checksum for the object when you use the SDK. This header will not provide any additional functionality if you don't use the SDK. When you send this header, there must be a corresponding x-amz-checksum or x-amz-trailer header sent. Otherwise, Amazon S3 fails the request with the HTTP status code 400 Bad Request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

If you provide an individual checksum, Amazon S3 ignores any provided ChecksumAlgorithm parameter.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 | XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in XML format.

Retention

Root level tag for the Retention parameters.

Required: Yes

EventHold

The event hold status for the object. Set to ON to enable an event hold or OFF to disable it.

Type: String

Valid Values: ON | OFF

Required: No

EventHoldDuration

The event hold duration for the object. Specifies how long the object remains protected after the event hold is released.

Type: [EventHoldDuration](#) data type

Required: No

Mode

Indicates the Retention mode for the specified object.

Type: String

Valid Values: GOVERNANCE | COMPLIANCE

Required: No

RetainUntilDate

The date on which this Object Lock Retention will expire.

Type: Timestamp

Required: No

Response Syntax

```
HTTP/1.1 200
x-amz-request-charged: RequestCharged
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutObjectTagging

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Sets the supplied tag-set to an object that already exists in a bucket. A tag is a key-value pair. For more information, see [Object Tagging](#).

You can associate tags with an object by sending a PUT request against the tagging subresource that is associated with the object. You can retrieve tags by sending a GET request. For more information, see [GetObjectTagging](#).

For tagging-related restrictions related to characters and encodings, see [Tag Restrictions](#). Note that Amazon S3 limits the maximum number of tags to 10 tags per object.

To use this operation, you must have permission to perform the `s3:PutObjectTagging` action. By default, the bucket owner has this permission and can grant this permission to others.

To put tags of any other version, use the `versionId` query parameter. You also need permission for the `s3:PutObjectVersionTagging` action.

`PutObjectTagging` has the following special errors. For more Amazon S3 errors see, [Error Responses](#).

- `InvalidTag` - The tag provided was not a valid tag. This error can occur if the tag did not pass input validation. For more information, see [Object Tagging](#).
- `MalformedXML` - The XML provided does not match the schema.
- `OperationAborted` - A conflicting conditional action is currently in progress against this resource. Please try again.
- `InternalServerError` - The service was unable to apply the provided tag to the object.

The following operations are related to `PutObjectTagging`:

- [GetObjectTagging](#)
- [DeleteObjectTagging](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /{Key+}?tagging&versionId=VersionId HTTP/1.1
Host: Bucket.s3.amazonaws.com
Content-MD5: ContentMD5
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-expected-bucket-owner: ExpectedBucketOwner
x-amz-request-payer: RequestPayer
<?xml version="1.0" encoding="UTF-8"?>
<Tagging xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <TagSet>
    <Tag>
      <Key>string</Key>
      <Value>string</Value>
    </Tag>
  </TagSet>
</Tagging>
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name containing the object.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form *AccessPointName-AccountId*.s3-accesspoint.*Region*.amazonaws.com. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of

the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

S3 on Outposts - When you use this action with S3 on Outposts, you must direct requests to the S3 on Outposts hostname. The S3 on Outposts hostname takes the form `AccessPointName-AccountId.outpostID.s3-outposts.Region.amazonaws.com`. When you use this action with S3 on Outposts, the destination bucket must be the Outposts access point ARN or the access point alias. For more information about S3 on Outposts, see [What is S3 on Outposts?](#) in the *Amazon S3 User Guide*.

Required: Yes

Content-MD5

The MD5 hash for the request body.

For requests made using the AWS Command Line Interface (CLI) or AWS SDKs, this field is calculated automatically.

Key

Name of the object key.

Length Constraints: Minimum length of 1.

Required: Yes

versionId

The versionId of the object that the tag-set will be added to.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-request-payer

Confirms that the requester knows that she or he will be charged for the tagging object request. Bucket owners need not specify this parameter in their requests.

Valid Values: requester

x-amz-sdk-checksum-algorithm

Indicates the algorithm used to create the checksum for the object when you use the SDK. This header will not provide any additional functionality if you don't use the SDK. When you send this header, there must be a corresponding `x-amz-checksum` or `x-amz-trailer` header sent. Otherwise, Amazon S3 fails the request with the HTTP status code 400 Bad Request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

If you provide an individual checksum, Amazon S3 ignores any provided `ChecksumAlgorithm` parameter.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 | XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in XML format.

Tagging

Root level tag for the Tagging parameters.

Required: Yes

TagSet

A collection for a set of tags

Type: Array of [Tag](#) data types

Required: Yes

Response Syntax

```
HTTP/1.1 200
x-amz-version-id: VersionId
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

x-amz-version-id

The versionId of the object the tag-set was added to.

Examples

Sample Request: Add tag set to an object

The following request adds a tag set to the existing object object-key in the examplebucket bucket.

```
PUT object-key?tagging HTTP/1.1
Host: examplebucket.s3.<Region>.amazonaws.com
Content-Length: length
Content-MD5: pUNXr/BjKK5G2UKEExample==
x-amz-date: 20160923T001956Z
Authorization: authorization string
<Tagging>
  <TagSet>
    <Tag>
      <Key>tag1</Key>
      <Value>val1</Value>
    </Tag>
    <Tag>
      <Key>tag2</Key>
      <Value>val2</Value>
    </Tag>
  </TagSet>
</Tagging>
```

Sample Response

This example illustrates one usage of PutObjectTagging.

```
HTTP/1.1 200 OK
x-amz-id-2: YgIPIfBiKa2bj0KMgUAdQkf3ShJT00pXUueF6QKo
x-amz-request-id: 236A8905248E5A01
```

Date: Fri, 23 Sep 2016 00:20:19 GMT

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutPublicAccessBlock

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Creates or modifies the `PublicAccessBlock` configuration for an Amazon S3 bucket. To use this operation, you must have the `s3:PutBucketPublicAccessBlock` permission. For more information about Amazon S3 permissions, see [Specifying Permissions in a Policy](#).

Important

When Amazon S3 evaluates the `PublicAccessBlock` configuration for a bucket or an object, it checks the `PublicAccessBlock` configuration for both the bucket (or the bucket that contains the object) and the bucket owner's account. Account-level settings automatically inherit from organization-level policies when present. If the `PublicAccessBlock` configurations are different between the bucket and the account, Amazon S3 uses the most restrictive combination of the bucket-level and account-level settings.

For more information about when Amazon S3 considers a bucket or an object public, see [The Meaning of "Public"](#).

The following operations are related to `PutPublicAccessBlock`:

- [GetPublicAccessBlock](#)
- [DeletePublicAccessBlock](#)
- [GetBucketPolicyStatus](#)
- [Using Amazon S3 Block Public Access](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /?publicAccessBlock HTTP/1.1
Host: Bucket.s3.amazonaws.com
Content-MD5: ContentMD5
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<PublicAccessBlockConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <BlockPublicAcls>boolean</BlockPublicAcls>
  <IgnorePublicAcls>boolean</IgnorePublicAcls>
  <BlockPublicPolicy>boolean</BlockPublicPolicy>
  <RestrictPublicBuckets>boolean</RestrictPublicBuckets>
</PublicAccessBlockConfiguration>
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the Amazon S3 bucket whose `PublicAccessBlock` configuration you want to set.

Required: Yes

Content-MD5

The MD5 hash of the `PutPublicAccessBlock` request body.

For requests made using the AWS Command Line Interface (CLI) or AWS SDKs, this field is calculated automatically.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-sdk-checksum-algorithm

Indicates the algorithm used to create the checksum for the object when you use the SDK. This header will not provide any additional functionality if you don't use the SDK. When you send this header, there must be a corresponding `x-amz-checksum` or `x-amz-trailer` header sent. Otherwise, Amazon S3 fails the request with the HTTP status code 400 Bad Request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

If you provide an individual checksum, Amazon S3 ignores any provided `ChecksumAlgorithm` parameter.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 | XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in XML format.

PublicAccessBlockConfiguration

Root level tag for the `PublicAccessBlockConfiguration` parameters.

Required: Yes

BlockPublicAcls

Specifies whether Amazon S3 should block public access control lists (ACLs) for this bucket and objects in this bucket. Setting this element to `TRUE` causes the following behavior:

- PUT Bucket ACL and PUT Object ACL calls fail if the specified ACL is public.
- PUT Object calls fail if the request includes a public ACL.
- PUT Bucket calls fail if the request includes a public ACL.

Enabling this setting doesn't affect existing policies or ACLs.

Type: Boolean

Required: No

BlockPublicPolicy

Specifies whether Amazon S3 should block public bucket policies for this bucket. Setting this element to TRUE causes Amazon S3 to reject calls to PUT Bucket policy if the specified bucket policy allows public access.

Enabling this setting doesn't affect existing bucket policies.

Type: Boolean

Required: No

IgnorePublicAcls

Specifies whether Amazon S3 should ignore public ACLs for this bucket and objects in this bucket. Setting this element to TRUE causes Amazon S3 to ignore all public ACLs on this bucket and objects in this bucket.

Enabling this setting doesn't affect the persistence of any existing ACLs and doesn't prevent new public ACLs from being set.

Type: Boolean

Required: No

RestrictPublicBuckets

Specifies whether Amazon S3 should restrict public bucket policies for this bucket. Setting this element to TRUE restricts access to this bucket to only AWS service principals and authorized users within this account if the bucket has a public policy.

Enabling this setting doesn't affect previously stored bucket policies, except that public and cross-account access within any public bucket policy, including non-public delegation to specific accounts, is blocked.

Type: Boolean

Required: No

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

First Sample Request

The following request puts a bucket `PublicAccessBlock` configuration that rejects public ACLs.

```
PUT /?publicAccessBlock HTTP/1.1
Host: <bucket-name>.s3.<Region>.amazonaws.com
x-amz-date: <Thu, 15 Nov 2016 00:17:21 GMT>
Authorization: <signatureValue>

<?xml version="1.0" encoding="UTF-8"?>
<PublicAccessBlockConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <BlockPublicAcls>TRUE</BlockPublicAcls>
  <IgnorePublicAcls>FALSE</IgnorePublicAcls>
  <BlockPublicPolicy>FALSE</BlockPublicPolicy>
  <RestrictPublicBuckets>FALSE</RestrictPublicBuckets>
</PublicAccessBlockConfiguration>
```

First Sample Response

This example illustrates one usage of `PutPublicAccessBlock`.

```
HTTP/1.1 200 OK
x-amz-id-2: ITnGT1y4REXAMPLEPi4hk1TXouTf0hccUjo0iCPEXAMPLEutBj3M7fPGLW02SEWp
x-amz-request-id: 51991EXAMPLE5321
Date: Thu, 15 Nov 2016 00:17:22 GMT
Server: AmazonS3
Content-Length: 0
```

Second Sample Request

The following request puts a bucket `PublicAccessBlock` configuration that ignores public ACLs and restricts access to public buckets.

```
PUT /?publicAccessBlock HTTP/1.1
Host: <bucket-name>.s3.<Region>.amazonaws.com
x-amz-date: <Thu, 15 Nov 2016 00:17:21 GMT>
Authorization: <signatureValue>

<?xml version="1.0" encoding="UTF-8"?>
<PublicAccessBlockConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <BlockPublicAcls>FALSE</BlockPublicAcls>
  <IgnorePublicAcls>TRUE</IgnorePublicAcls>
  <BlockPublicPolicy>FALSE</BlockPublicPolicy>
  <RestrictPublicBuckets>TRUE</RestrictPublicBuckets>
</PublicAccessBlockConfiguration>
```

Second Sample Response

This example illustrates one usage of PutPublicAccessBlock.

```
HTTP/1.1 200 OK
x-amz-id-2: ITnGT1y4REXAMPLEPi4hk1TXouTf0hccUjo0iCPEXAMPLEutBj3M7fPG1W02SEWp
x-amz-request-id: 51991EXAMPLE5321
Date: Thu, 15 Nov 2016 00:17:22 GMT
Server: AmazonS3
Content-Length: 0
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)

- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

RenameObject

Service: Amazon S3

Renames an existing object in a directory bucket that uses the S3 Express One Zone storage class. You can use `RenameObject` by specifying an existing object's name as the source and the new name of the object as the destination within the same directory bucket.

Note

`RenameObject` is only supported for objects stored in the S3 Express One Zone storage class.

To prevent overwriting an object, you can use the `If-None-Match` conditional header.

- **If-None-Match** - Renames the object only if an object with the specified name does not already exist in the directory bucket. If you don't want to overwrite an existing object, you can add the `If-None-Match` conditional header with the value `'*'` in the `RenameObject` request. Amazon S3 then returns a `412 Precondition Failed` error if the object with the specified name already exists. For more information, see [RFC 7232](#).

Permissions

To grant access to the `RenameObject` operation on a directory bucket, we recommend that you use the `CreateSession` operation for session-based authorization. Specifically, you grant the `s3express:CreateSession` permission to the directory bucket in a bucket policy or an IAM identity-based policy. Then, you make the `CreateSession` API call on the directory bucket to obtain a session token. With the session token in your request header, you can make API requests to this operation. After the session token expires, you make another `CreateSession` API call to generate a new session token for use. The AWS CLI and SDKs will create and manage your session including refreshing the session token automatically to avoid service interruptions when a session expires. In your bucket policy, you can specify the `s3express:SessionMode` condition key to control who can create a `ReadWrite` or `ReadOnly` session. A `ReadWrite` session is required for executing all the Zonal endpoint API operations, including `RenameObject`. For more information about authorization, see [CreateSession](#). To learn more about Zonal endpoint API operations, see [Authorizing Zonal endpoint API operations with CreateSession](#) in the *Amazon S3 User Guide*.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `Bucket-name.s3express-zone-id.region-code.amazonaws.com`.

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /{Key+}?renameObject HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-rename-source: RenameSource
If-Match: DestinationIfMatch
If-None-Match: DestinationIfNoneMatch
If-Modified-Since: DestinationIfModifiedSince
If-Unmodified-Since: DestinationIfUnmodifiedSince
x-amz-rename-source-if-match: SourceIfMatch
x-amz-rename-source-if-none-match: SourceIfNoneMatch
x-amz-rename-source-if-modified-since: SourceIfModifiedSince
x-amz-rename-source-if-unmodified-since: SourceIfUnmodifiedSince
x-amz-client-token: ClientToken
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name of the directory bucket containing the object.

You must use virtual-hosted-style requests in the format `Bucket-name.s3express-zone-id.region-code.amazonaws.com`. Path-style requests are not supported. Directory bucket names must be unique in the chosen Availability Zone. Bucket names must follow the format `bucket-base-name--zone-id--x-s3` (for example, `amzn-s3-demo-bucket--usw2-`

az1--x-s3). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*.

Required: Yes

If-Match

Renames the object only if the ETag (entity tag) value provided during the operation matches the ETag of the object in S3. The If-Match header field makes the request method conditional on ETags. If the ETag values do not match, the operation returns a 412 Precondition Failed error.

Expects the ETag value as a string.

If-Modified-Since

Renames the object if the destination exists and if it has been modified since the specified time.

If-None-Match

Renames the object only if the destination does not already exist in the specified directory bucket. If the object does exist when you send a request with If-None-Match:*, the S3 API will return a 412 Precondition Failed error, preventing an overwrite. The If-None-Match header prevents overwrites of existing data by validating that there's not an object with the same key name already in your directory bucket.

Expects the * character (asterisk).

If-Unmodified-Since

Renames the object if it hasn't been modified since the specified time.

Key

Key name of the object to rename.

Length Constraints: Minimum length of 1.

Required: Yes

x-amz-client-token

A unique string with a max of 64 ASCII characters in the ASCII range of 33 - 126.

Note

RenameObject supports idempotency using a client token. To make an idempotent API request using RenameObject, specify a client token in the request. You should not reuse the same client token for other API requests. If you retry a request that completed successfully using the same client token and the same parameters, the retry succeeds without performing any further actions. If you retry a successful request using the same client token, but one or more of the parameters are different, the retry fails and an IdempotentParameterMismatch error is returned.

x-amz-rename-source

Specifies the source for the rename operation. The value must be URL encoded.

Pattern: `\/?\.\+\/\.\+`

Required: Yes

x-amz-rename-source-if-match

Renames the object if the source exists and if its entity tag (ETag) matches the specified ETag.

x-amz-rename-source-if-modified-since

Renames the object if the source exists and if it has been modified since the specified time.

x-amz-rename-source-if-none-match

Renames the object if the source exists and if its entity tag (ETag) is different than the specified ETag. If an asterisk (*) character is provided, the operation will fail and return a 412 Precondition Failed error.

x-amz-rename-source-if-unmodified-since

Renames the object if the source exists and hasn't been modified since the specified time.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

IdempotencyParameterMismatch

Parameters on this idempotent request are inconsistent with parameters used in previous request(s).

For a list of error codes and more information on Amazon S3 errors, see [Error codes](#).

Note

Idempotency ensures that an API request completes no more than one time. With an idempotent request, if the original request completes successfully, any subsequent retries complete successfully without performing any further actions.

HTTP Status Code: 400

Examples

Sample request: Renaming an object in a directory bucket

The following example request illustrates renaming the *srcfilename.txt* object in a directory bucket.

```
PUT /dstfilename.txt HTTP/1.1
Host: amzn-s3-demo-bucket--zone-id--x-s3.s3express.amazonaws.com
If-Match: 9b2cf535f27731c974343645a3985328
x-amz-rename-source: /srcfilename.txt
```

Sample request: Renaming an object in a directory bucket if the destination doesn't exist

The following example request illustrates renaming the *srcfilename.txt* object to *dstfilename.txt* if an object named *dstfilename.txt* doesn't already exist in the directory bucket.

```
PUT /dstfilename.txt HTTP/1.1
Host: amzn-s3-demo-bucket--zone-id--x-s3.s3express.amazonaws.com
If-None-Match: *
x-amz-rename-source: /srcfilename.txt
```

Sample request: Renaming an object in a directory bucket if the destination exists with a specific ETag

The following example request illustrates renaming the *srcfilename.txt* object only if the destination exists and has the ETag: *9b2cf535f27731c974343645a3985328*.

```
PUT /dstfilename.txt HTTP/1.1
Host: amzn-s3-demo-bucket--zone-id--x-s3.s3express.amazonaws.com
If-Match: 9b2cf535f27731c974343645a3985328
x-amz-rename-source: /srcfilename.txt
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

RestoreObject

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Restores an archived copy of an object back into Amazon S3

This functionality is not supported for Amazon S3 on Outposts.

This action performs the following types of requests:

- `restore an archive` - Restore an archived object

For more information about the S3 structure in the request body, see the following:

- [PutObject](#)
- [Managing Access with ACLs](#) in the *Amazon S3 User Guide*
- [Protecting Data Using Server-Side Encryption](#) in the *Amazon S3 User Guide*

Permissions

To use this operation, you must have permissions to perform the `s3:RestoreObject` action. The bucket owner has this permission by default and can grant this permission to others. For more information about permissions, see [Permissions Related to Bucket Subresource Operations](#) and [Managing Access Permissions to Your Amazon S3 Resources](#) in the *Amazon S3 User Guide*.

Restoring objects

Objects that you archive to the S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive storage class, and S3 Intelligent-Tiering Archive or S3 Intelligent-Tiering Deep Archive tiers, are not accessible in real time. For objects in the S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive storage classes, you must first initiate a restore request, and then wait until a temporary copy of the object is available. If you want a permanent copy of the object, create a copy of it in the Amazon S3 Standard storage class in your S3 bucket. To access an archived object, you must restore the object for the duration (number of days) that you specify. For objects in the Archive

Access or Deep Archive Access tiers of S3 Intelligent-Tiering, you must first initiate a restore request, and then wait until the object is moved into the Frequent Access tier.

To restore a specific object version, you can provide a version ID. If you don't provide a version ID, Amazon S3 restores the current version.

When restoring an archived object, you can specify one of the following data access tier options in the `Tier` element of the request body:

- **Expedited** - Expedited retrievals allow you to quickly access your data stored in the S3 Glacier Flexible Retrieval storage class or S3 Intelligent-Tiering Archive tier when occasional urgent requests for restoring archives are required. For all but the largest archived objects (250 MB+), data accessed using Expedited retrievals is typically made available within 1–5 minutes. Provisioned capacity ensures that retrieval capacity for Expedited retrievals is available when you need it. Expedited retrievals and provisioned capacity are not available for objects stored in the S3 Glacier Deep Archive storage class or S3 Intelligent-Tiering Deep Archive tier.
- **Standard** - Standard retrievals allow you to access any of your archived objects within several hours. This is the default option for retrieval requests that do not specify the retrieval option. Standard retrievals typically finish within 3–5 hours for objects stored in the S3 Glacier Flexible Retrieval storage class or S3 Intelligent-Tiering Archive tier. They typically finish within 12 hours for objects stored in the S3 Glacier Deep Archive storage class or S3 Intelligent-Tiering Deep Archive tier. Standard retrievals are free for objects stored in S3 Intelligent-Tiering.
- **Bulk** - Bulk retrievals free for objects stored in the S3 Glacier Flexible Retrieval and S3 Intelligent-Tiering storage classes, enabling you to retrieve large amounts, even petabytes, of data at no cost. Bulk retrievals typically finish within 5–12 hours for objects stored in the S3 Glacier Flexible Retrieval storage class or S3 Intelligent-Tiering Archive tier. Bulk retrievals are also the lowest-cost retrieval option when restoring objects from S3 Glacier Deep Archive. They typically finish within 48 hours for objects stored in the S3 Glacier Deep Archive storage class or S3 Intelligent-Tiering Deep Archive tier.

For more information about archive retrieval options and provisioned capacity for Expedited data access, see [Restoring Archived Objects](#) in the *Amazon S3 User Guide*.

You can use Amazon S3 restore speed upgrade to change the restore speed to a faster speed while it is in progress. For more information, see [Upgrading the speed of an in-progress restore](#) in the *Amazon S3 User Guide*.

To get the status of object restoration, you can send a HEAD request. Operations return the `x-amz-restore` header, which provides information about the restoration status, in the response. You can use Amazon S3 event notifications to notify you when a restore is initiated or completed. For more information, see [Configuring Amazon S3 Event Notifications](#) in the *Amazon S3 User Guide*.

After restoring an archived object, you can update the restoration period by reissuing the request with a new period. Amazon S3 updates the restoration period relative to the current time and charges only for the request—there are no data transfer charges. You cannot update the restoration period when Amazon S3 is actively processing your current restore request for the object.

If your bucket has a lifecycle configuration with a rule that includes an expiration action, the object expiration overrides the life span that you specify in a restore request. For example, if you restore an object copy for 10 days, but the object is scheduled to expire in 3 days, Amazon S3 deletes the object in 3 days. For more information about lifecycle configuration, see [PutBucketLifecycleConfiguration](#) and [Object Lifecycle Management](#) in *Amazon S3 User Guide*.

Responses

A successful action returns either the 200 OK or 202 Accepted status code.

- If the object is not previously restored, then Amazon S3 returns 202 Accepted in the response.
- If the object is previously restored, Amazon S3 returns 200 OK in the response.
- Special errors:
 - *Code: RestoreAlreadyInProgress*
 - *Cause: Object restore is already in progress.*
 - *HTTP Status Code: 409 Conflict*
 - *SOAP Fault Code Prefix: Client*
- *Code: GlacierExpeditedRetrievalNotAvailable*
 - *Cause: expedited retrievals are currently not available. Try again later. (Returned if there is insufficient capacity to process the Expedited request. This error applies only to Expedited retrievals and not to S3 Standard or Bulk retrievals.)*
 - *HTTP Status Code: 503*
 - *SOAP Fault Code Prefix: N/A*

The following operations are related to `RestoreObject`:

- [PutBucketLifecycleConfiguration](#)
- [GetBucketNotificationConfiguration](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
POST /{Key+}?restore&versionId=VersionId HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-request-payer: RequestPayer
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<RestoreRequest xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Days>integer</Days>
  <GlacierJobParameters>
    <Tier>string</Tier>
  </GlacierJobParameters>
  <Type>string</Type>
  <Tier>string</Tier>
  <Description>string</Description>
  <SelectParameters>
    <Expression>string</Expression>
    <ExpressionType>string</ExpressionType>
    <InputSerialization>
      <CompressionType>string</CompressionType>
      <CSV>
        <AllowQuotedRecordDelimiter>boolean</AllowQuotedRecordDelimiter>
        <Comments>string</Comments>
        <FieldDelimiter>string</FieldDelimiter>
        <FileHeaderInfo>string</FileHeaderInfo>
        <QuoteCharacter>string</QuoteCharacter>
        <QuoteEscapeCharacter>string</QuoteEscapeCharacter>
        <RecordDelimiter>string</RecordDelimiter>
      </CSV>
    </InputSerialization>
  </SelectParameters>
</RestoreRequest>
```

```

</CSV>
<JSON>
  <Type>string</Type>
</JSON>
<Parquet>
</Parquet>
</InputSerialization>
<OutputSerialization>
  <CSV>
    <FieldDelimiter>string</FieldDelimiter>
    <QuoteCharacter>string</QuoteCharacter>
    <QuoteEscapeCharacter>string</QuoteEscapeCharacter>
    <QuoteFields>string</QuoteFields>
    <RecordDelimiter>string</RecordDelimiter>
  </CSV>
  <JSON>
    <RecordDelimiter>string</RecordDelimiter>
  </JSON>
</OutputSerialization>
</SelectParameters>
<OutputLocation>
  <S3>
    <AccessControlList>
      <Grant>
        <Grantee>
          <DisplayName>string</DisplayName>
          <EmailAddress>string</EmailAddress>
          <ID>string</ID>
          <xsi:type>string</xsi:type>
          <URI>string</URI>
        </Grantee>
        <Permission>string</Permission>
      </Grant>
    </AccessControlList>
    <BucketName>string</BucketName>
    <CannedACL>string</CannedACL>
    <Encryption>
      <EncryptionType>string</EncryptionType>
      <KMSContext>string</KMSContext>
      <KMSKeyId>string</KMSKeyId>
    </Encryption>
    <Prefix>string</Prefix>
    <StorageClass>string</StorageClass>
    <Tagging>

```

```

    <TagSet>
      <Tag>
        <Key>string</Key>
        <Value>string</Value>
      </Tag>
    </TagSet>
  </Tagging>
  <UserMetadata>
    <MetadataEntry>
      <Name>string</Name>
      <Value>string</Value>
    </MetadataEntry>
  </UserMetadata>
</S3>
</OutputLocation>
</RestoreRequest>

```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name containing the object to restore.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form *AccessPointName-AccountId.s3-accesspoint.Region.amazonaws.com*. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

S3 on Outposts - When you use this action with S3 on Outposts, you must direct requests to the S3 on Outposts hostname. The S3 on Outposts hostname takes the form *AccessPointName-AccountId.outpostID.s3-outposts.Region.amazonaws.com*. When you use this action with S3 on Outposts, the destination bucket must be the Outposts access point ARN or the access point alias. For more information about S3 on Outposts, see [What is S3 on Outposts?](#) in the *Amazon S3 User Guide*.

Required: Yes

Key

Object key for which the action was initiated.

Length Constraints: Minimum length of 1.

Required: Yes

versionId

VersionId used to reference a specific version of the object.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

x-amz-sdk-checksum-algorithm

Indicates the algorithm used to create the checksum for the object when you use the SDK. This header will not provide any additional functionality if you don't use the SDK. When you send this header, there must be a corresponding x-amz-checksum or x-amz-trailer header sent. Otherwise, Amazon S3 fails the request with the HTTP status code 400 Bad Request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

If you provide an individual checksum, Amazon S3 ignores any provided `ChecksumAlgorithm` parameter.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 | XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in XML format.

RestoreRequest

Root level tag for the `RestoreRequest` parameters.

Required: Yes

Days

Lifetime of the active copy in days. Do not use with restores that specify `OutputLocation`.

The `Days` element is required for regular restores, and must not be provided for select requests.

Type: Integer

Required: No

Description

The optional description for the job.

Type: String

Required: No

GlacierJobParameters

S3 Glacier related parameters pertaining to this job. Do not use with restores that specify `OutputLocation`.

Type: [GlacierJobParameters](#) data type

Required: No

OutputLocation

Describes the location where the restore job's output is stored.

Type: [OutputLocation](#) data type

Required: No

SelectParameters

Important

Amazon S3 Select is no longer available to new customers. Existing customers of Amazon S3 Select can continue to use the feature as usual. [Learn more](#)

Describes the parameters for Select job types.

Type: [SelectParameters](#) data type

Required: No

Tier

Retrieval tier at which the restore will be processed.

Type: String

Valid Values: Standard | Bulk | Expedited

Required: No

Type

Important

Amazon S3 Select is no longer available to new customers. Existing customers of Amazon S3 Select can continue to use the feature as usual. [Learn more](#)

Type of restore request.

Type: String

Valid Values: SELECT

Required: No

Response Syntax

```
HTTP/1.1 200
x-amz-request-charged: RequestCharged
x-amz-restore-output-path: RestoreOutputPath
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

x-amz-restore-output-path

Indicates the path in the provided S3 output location where Select results will be restored to.

Errors

ObjectAlreadyInActiveTierError

This action is not allowed against this storage tier.

HTTP Status Code: 403

Examples

Example: Restore an object for 2 days using the expedited retrieval option

The following restore request restores a copy of the `photo1.jpg` object from S3 Glacier for a period of two days using the expedited retrieval option.

```
POST /photo1.jpg?restore HTTP/1.1
Host: examplebucket.dummy value
Date: Mon, 22 Oct 2012 01:49:52 GMT
Authorization: authorization string
Content-Length: content length
<RestoreRequest>
  <Days>2</Days>
  <GlacierJobParameters>
    <Tier>Standard</Tier>
  </GlacierJobParameters>
</RestoreRequest>
```

Sample response

If the `examplebucket` does not have a restored copy of the object, Amazon S3 returns the following `202 Accepted` response.

Note

If a copy of the object is already restored, Amazon S3 returns a `200 OK` response, and updates only the restored copy's expiry time.

```
HTTP/1.1 202 Accepted
x-amz-id-2: GFihv3y6+kE7KG11GEkQhU7/2/cHR3Yb2fCb2S04nxI423Dqwg2XiQ0B/
UZ1zYQvPiBlZNRcovw=
x-amz-request-id: 9F341CD3C4BA79E0
Date: Sat, 20 Oct 2012 23:54:05 GMT
Content-Length: 0
Server: AmazonS3
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

SelectObjectContent

Service: Amazon S3

Note

This operation is not supported for directory buckets.

This action filters the contents of an Amazon S3 object based on a simple structured query language (SQL) statement. In the request, along with the SQL expression, you must also specify a data serialization format (JSON, CSV, or Apache Parquet) of the object. Amazon S3 uses this format to parse object data into records, and returns only records that match the specified SQL expression. You must also specify the data serialization format for the response.

This functionality is not supported for Amazon S3 on Outposts.

For more information about Amazon S3 Select, see [Selecting Content from Objects](#) and [SELECT Command](#) in the *Amazon S3 User Guide*.

Permissions

You must have the `s3:GetObject` permission for this operation. Amazon S3 Select does not support anonymous access. For more information about permissions, see [Specifying Permissions in a Policy](#) in the *Amazon S3 User Guide*.

Object Data Formats

You can use Amazon S3 Select to query objects that have the following format properties:

- *CSV, JSON, and Parquet* - Objects must be in CSV, JSON, or Parquet format.
- *UTF-8* - UTF-8 is the only encoding type Amazon S3 Select supports.
- *GZIP or BZIP2* - CSV and JSON files can be compressed using GZIP or BZIP2. GZIP and BZIP2 are the only compression formats that Amazon S3 Select supports for CSV and JSON files. Amazon S3 Select supports columnar compression for Parquet using GZIP or Snappy. Amazon S3 Select does not support whole-object compression for Parquet objects.
- *Server-side encryption* - Amazon S3 Select supports querying objects that are protected with server-side encryption.

For objects that are encrypted with customer-provided encryption keys (SSE-C), you must use HTTPS, and you must use the headers that are documented in the [GetObject](#). For more

information about SSE-C, see [Server-Side Encryption \(Using Customer-Provided Encryption Keys\)](#) in the *Amazon S3 User Guide*.

For objects that are encrypted with Amazon S3 managed keys (SSE-S3) and AWS KMS keys (SSE-KMS), server-side encryption is handled transparently, so you don't need to specify anything. For more information about server-side encryption, including SSE-S3 and SSE-KMS, see [Protecting Data Using Server-Side Encryption](#) in the *Amazon S3 User Guide*.

Working with the Response Body

Given the response size is unknown, Amazon S3 Select streams the response as a series of messages and includes a `Transfer-Encoding` header with `chunked` as its value in the response. For more information, see [Appendix: SelectObjectContent Response](#).

GetObject Support

The `SelectObjectContent` action does not support the following `GetObject` functionality. For more information, see [GetObject](#).

- **Range:** Although you can specify a scan range for an Amazon S3 Select request (see [SelectObjectContentRequest - ScanRange](#) in the request parameters), you cannot specify the range of bytes of an object to return.
- The `GLACIER`, `DEEP_ARCHIVE`, and `REDUCED_REDUNDANCY` storage classes, or the `ARCHIVE_ACCESS` and `DEEP_ARCHIVE_ACCESS` access tiers of the `INTELLIGENT_TIERING` storage class: You cannot query objects in the `GLACIER`, `DEEP_ARCHIVE`, or `REDUCED_REDUNDANCY` storage classes, nor objects in the `ARCHIVE_ACCESS` or `DEEP_ARCHIVE_ACCESS` access tiers of the `INTELLIGENT_TIERING` storage class. For more information about storage classes, see [Using Amazon S3 storage classes](#) in the *Amazon S3 User Guide*.

Special Errors

For a list of special errors for this operation, see [List of SELECT Object Content Error Codes](#)

The following operations are related to `SelectObjectContent`:

- [GetObject](#)
- [GetBucketLifecycleConfiguration](#)
- [PutBucketLifecycleConfiguration](#)

⚠ Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
POST /{Key+}?select&select-type=2 HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-server-side-encryption-customer-algorithm: SSECustomerAlgorithm
x-amz-server-side-encryption-customer-key: SSECustomerKey
x-amz-server-side-encryption-customer-key-MD5: SSECustomerKeyMD5
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<SelectObjectContentRequest xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Expression>string</Expression>
  <ExpressionType>string</ExpressionType>
  <RequestProgress>
    <Enabled>boolean</Enabled>
  </RequestProgress>
  <InputSerialization>
    <CompressionType>string</CompressionType>
    <CSV>
      <AllowQuotedRecordDelimiter>boolean</AllowQuotedRecordDelimiter>
      <Comments>string</Comments>
      <FieldDelimiter>string</FieldDelimiter>
      <FileHeaderInfo>string</FileHeaderInfo>
      <QuoteCharacter>string</QuoteCharacter>
      <QuoteEscapeCharacter>string</QuoteEscapeCharacter>
      <RecordDelimiter>string</RecordDelimiter>
    </CSV>
    <JSON>
      <Type>string</Type>
    </JSON>
    <Parquet>
  </Parquet>
</InputSerialization>
<OutputSerialization>
  <CSV>
    <FieldDelimiter>string</FieldDelimiter>
```

```
<QuoteCharacter>string</QuoteCharacter>
<QuoteEscapeCharacter>string</QuoteEscapeCharacter>
<QuoteFields>string</QuoteFields>
<RecordDelimiter>string</RecordDelimiter>
</CSV>
<JSON>
  <RecordDelimiter>string</RecordDelimiter>
</JSON>
</OutputSerialization>
<ScanRange>
  <End>long</End>
  <Start>long</Start>
</ScanRange>
</SelectObjectContentRequest>
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The S3 bucket.

Required: Yes

Key

The object key.

Length Constraints: Minimum length of 1.

Required: Yes

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-server-side-encryption-customer-algorithm

The server-side encryption (SSE) algorithm used to encrypt the object. This parameter is needed only when the object was created using a checksum algorithm. For more information, see [Protecting data using SSE-C keys](#) in the *Amazon S3 User Guide*.

[x-amz-server-side-encryption-customer-key](#)

The server-side encryption (SSE) customer managed key. This parameter is needed only when the object was created using a checksum algorithm. For more information, see [Protecting data using SSE-C keys](#) in the *Amazon S3 User Guide*.

[x-amz-server-side-encryption-customer-key-MD5](#)

The MD5 server-side encryption (SSE) customer managed key. This parameter is needed only when the object was created using a checksum algorithm. For more information, see [Protecting data using SSE-C keys](#) in the *Amazon S3 User Guide*.

Request Body

The request accepts the following data in XML format.

[SelectObjectContentRequest](#)

Root level tag for the SelectObjectContentRequest parameters.

Required: Yes

[Expression](#)

The expression that is used to query the object.

Type: String

Required: Yes

[ExpressionType](#)

The type of the provided expression (for example, SQL).

Type: String

Valid Values: SQL

Required: Yes

[InputSerialization](#)

Describes the format of the data in the object that is being queried.

Type: [InputSerialization](#) data type

Required: Yes

OutputSerialization

Describes the format of the data that you want Amazon S3 to return in response.

Type: [OutputSerialization](#) data type

Required: Yes

RequestProgress

Specifies if periodic request progress information should be enabled.

Type: [RequestProgress](#) data type

Required: No

ScanRange

Specifies the byte range of the object to get the records from. A record is processed when its first byte is contained by the range. This parameter is optional, but when specified, it must not be empty. See RFC 2616, Section 14.35.1 about how to specify the start and end of the range.

ScanRange may be used in the following ways:

- `<scanrange><start>50</start><end>100</end></scanrange>` - process only the records starting between the bytes 50 and 100 (inclusive, counting from zero)
- `<scanrange><start>50</start></scanrange>` - process only the records starting after the byte 50
- `<scanrange><end>50</end></scanrange>` - process only the records within the last 50 bytes of the file.

Type: [ScanRange](#) data type

Required: No

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<Payload>
  <Records>
```

```

    <Payload>blob</Payload>
  </Records>
  <Stats>
    <Details>
      <BytesProcessed>long</BytesProcessed>
      <BytesReturned>long</BytesReturned>
      <BytesScanned>long</BytesScanned>
    </Details>
  </Stats>
  <Progress>
    <Details>
      <BytesProcessed>long</BytesProcessed>
      <BytesReturned>long</BytesReturned>
      <BytesScanned>long</BytesScanned>
    </Details>
  </Progress>
  <Cont>
  </Cont>
  <End>
  </End>
</Payload>

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

Payload

Root level tag for the Payload parameters.

Required: Yes

Cont

The Continuation Event.

Type: [ContinuationEvent](#) data type

End

The End Event.

Type: [EndEvent](#) data type

Progress

The Progress Event.

Type: [ProgressEvent](#) data type

Records

The Records Event.

Type: [RecordsEvent](#) data type

Stats

The Stats Event.

Type: [StatsEvent](#) data type

Examples

Example 1: CSV object

The following select request retrieves all records from an object with data stored in CSV format. The OutputSerialization element directs Amazon S3 to return results in CSV.

You can try different queries in the Expression element:

- Assuming that you are not using column headers, you can identify columns using positional headers:

```
SELECT s._1, s._2 FROM S3Object s WHERE s._3 > 100
```

- If you have column headers and you set the FileHeaderInfo to Use, you can identify columns by name in the expression:

```
SELECT s.Id, s.FirstName, s.SSN FROM S3Object s
```

- You can specify functions in the SQL expression:

```
SELECT count(*) FROM S3Object s WHERE s._1 < 1
```

```
POST /exampleobject.csv?select&select-type=2 HTTP/1.1
```

```

Host: examplebucket.s3.<Region>.amazonaws.com
Date: Tue, 17 Oct 2017 01:49:52 GMT
Authorization: authorization string
Content-Length: content length

<?xml version="1.0" encoding="UTF-8"?>
<SelectRequest>
  <Expression>Select * from S3Object</Expression>
  <ExpressionType>SQL</ExpressionType>
  <InputSerialization>
    <CompressionType>GZIP</CompressionType>
    <CSV>
      <FileHeaderInfo>IGNORE</FileHeaderInfo>
      <RecordDelimiter>\n</RecordDelimiter>
      <FieldDelimiter>,</FieldDelimiter>
      <QuoteCharacter>"</QuoteCharacter>
      <QuoteEscapeCharacter>"</QuoteEscapeCharacter>
      <Comments>#</Comments>
    </CSV>
  </InputSerialization>
  <OutputSerialization>
    <CSV>
      <QuoteFields>ASNEEDED</QuoteFields>
      <RecordDelimiter>\n</RecordDelimiter>
      <FieldDelimiter>,</FieldDelimiter>
      <QuoteCharacter>"</QuoteCharacter>
      <QuoteEscapeCharacter>"</QuoteEscapeCharacter>
    </CSV>
  </OutputSerialization>
</SelectRequest>

```

Example

The following is a sample response.

```

HTTP/1.1 200 OK
x-amz-id-2: GFihv3y6+kE7KG11GEkQhU7/2/cHR3Yb2fCb2S04nxI423Dqwg2XiQ0B/
UZ1zYQvPiBlZNRcovw=
x-amz-request-id: 9F341CD3C4BA79E0
Date: Tue, 17 Oct 2017 23:54:05 GMT

```

A series of messages

Example 2: JSON object

The following select request retrieves all records from an object with data stored in JSON format. The OutputSerialization directs Amazon S3 to return results in CSV.

You can try different queries in the Expression element:

- You can filter by string comparison using record keys:

```
SELECT s.country, s.city from S3Object s where s.city = 'Seattle'
```

- You can specify functions in the SQL expression:

```
SELECT count(*) FROM S3Object s
```

```
POST /exampleobject.json?select&select-type=2 HTTP/1.1
```

```
Host: examplebucket.s3.<Region>.amazonaws.com
```

```
Date: Tue, 17 Oct 2017 01:49:52 GMT
```

```
Authorization: authorization string
```

```
Content-Length: content length
```

```
<?xml version="1.0" encoding="UTF-8"?>
```

```
<SelectRequest>
```

```
  <Expression>Select * from S3Object</Expression>
```

```
  <ExpressionType>SQL</ExpressionType>
```

```
  <InputSerialization>
```

```
    <CompressionType>GZIP</CompressionType>
```

```
    <JSON>
```

```
      <Type>DOCUMENT</Type>
```

```
    </JSON>
```

```
  </InputSerialization>
```

```
  <OutputSerialization>
```

```
    <CSV>
```

```
      <QuoteFields>ASNEEDED</QuoteFields>
```

```
      <RecordDelimiter>\n</RecordDelimiter>
```

```
      <FieldDelimiter>,</FieldDelimiter>
```

```
      <QuoteCharacter>"</QuoteCharacter>
```

```
      <QuoteEscapeCharacter>"</QuoteEscapeCharacter>
```

```
    </CSV>
```

```
  </OutputSerialization>
```

```
</SelectRequest>
```

Example

The following is a sample response.

```
HTTP/1.1 200 OK
x-amz-id-2: GFihv3y6+kE7KG11GEkQhU7/2/CHR3Yb2fCb2S04nxI423Dqwg2XiQ0B/
UZlzyQvPiBlZNRcovw=
x-amz-request-id: 9F341CD3C4BA79E0
Date: Tue, 17 Oct 2017 23:54:05 GMT

A series of messages
```

Example 3: Parquet object

- The `InputSerialization` element describes the format of the data in the object that is being queried. It must specify CSV, JSON, or Parquet.
- The `OutputSerialization` element describes the format of the data that you want Amazon S3 to return in response to the query. It must specify CSV, JSON. Amazon S3 doesn't support outputting data in the Parquet format.
- The format of the `InputSerialization` doesn't need to match the format of the `OutputSerialization`. So, for example, you can specify JSON in the `InputSerialization` and CSV in the `OutputSerialization`.

```
POST /exampleobject.parquet?select&select-type=2 HTTP/1.1
Host: examplebucket.s3.<Region>.amazonaws.com
Date: Tue, 17 Oct 2017 01:49:52 GMT
Authorization: authorization string
Content-Length: content length

<?xml version="1.0" encoding="UTF-8"?>
<SelectRequest>
  <Expression>Select * from S3object</Expression>
  <ExpressionType>SQL</ExpressionType>
  <InputSerialization>
```

```
<CompressionType>NONE</CompressionType>
<Parquet>
</Parquet>
</InputSerialization>
<OutputSerialization>
  <CSV>
    <QuoteFields>ASNEEDED</QuoteFields>
    <RecordDelimiter>\n</RecordDelimiter>
    <FieldDelimiter>,</FieldDelimiter>
    <QuoteCharacter>"</QuoteCharacter>
    <QuoteEscapeCharacter>"</QuoteEscapeCharacter>
  </CSV>
</OutputSerialization>
</SelectRequest>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

UpdateBucketMetadataAnnotationTableConfiguration

Service: Amazon S3

Updates the annotation table configuration for an Amazon S3 bucket's metadata configuration. Use this operation to enable or disable the annotation table, or to update its associated IAM role.

An annotation table is a queryable Iceberg table that contains records of all annotations attached to objects in the bucket. To use this operation, the bucket must have an existing Amazon S3 Metadata configuration.

To use this operation, you must have the `s3:UpdateBucketMetadataAnnotationTableConfiguration` permission. If you are specifying or changing the IAM role, you must also have `iam:PassRole` permission for the role.

The IAM role must have a trust policy that allows the Amazon S3 metadata service to assume it, and a permissions policy that grants the actions needed to read annotations from your bucket. The following examples show a trust policy and a permissions policy that you can adapt for your bucket and account.

The following operations are related to `UpdateBucketMetadataAnnotationTableConfiguration`:

- [CreateBucketMetadataConfiguration](#)
- [GetBucketMetadataConfiguration](#)

Request Syntax

```
PUT /?metadataAnnotationTable HTTP/1.1
Host: Bucket.s3.amazonaws.com
Content-MD5: ContentMD5
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<AnnotationTableConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <ConfigurationState>string</ConfigurationState>
  <EncryptionConfiguration>
    <KmsKeyArn>string</KmsKeyArn>
    <SseAlgorithm>string</SseAlgorithm>
  </EncryptionConfiguration>
  <Role>string</Role>
```

```
</AnnotationTableConfiguration>
```

URI Request Parameters

The request uses the following URI parameters.

[Bucket](#)

The name of the bucket whose annotation table configuration to update.

Required: Yes

[Content-MD5](#)

Base64-encoded MD5 digest of the message body.

[x-amz-expected-bucket-owner](#)

The account ID of the expected bucket owner.

[x-amz-sdk-checksum-algorithm](#)

Checksum algorithm for the request payload.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 |
XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in XML format.

[AnnotationTableConfiguration](#)

Root level tag for the AnnotationTableConfiguration parameters.

Required: Yes

[ConfigurationState](#)

The new configuration state to apply.

Type: String

Valid Values: ENABLED | DISABLED

Required: Yes

EncryptionConfiguration

The encryption settings for an S3 Metadata journal table or inventory table configuration.

Type: [MetadataTableEncryptionConfiguration](#) data type

Required: No

Role

The new IAM role ARN to apply.

Type: String

Required: No

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Trust policy

This example illustrates one usage of `UpdateBucketMetadataAnnotationTableConfiguration`.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "metadata.s3.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
```

```

        "aws:SourceAccount": "123456789012"
      },
      "ArnLike": {
        "aws:SourceArn": "arn:aws:s3:::amzn-s3-demo-bucket"
      }
    }
  ]
}

```

Permissions policy

This example illustrates one usage of `UpdateBucketMetadataAnnotationTableConfiguration`.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "PermissionForGetAnnotation",
      "Effect": "Allow",
      "Action": [
        "s3:GetObjectAnnotation",
        "s3:GetObjectVersionAnnotation"
      ],
      "Resource": ["arn:aws:s3:::amzn-s3-demo-bucket/*"],
      "Condition": {
        "StringEquals": {
          "aws:ResourceAccount": "123456789012"
        }
      }
    },
    {
      "Sid": "PermissionsForListBucket",
      "Effect": "Allow",
      "Action": [
        "s3:ListBucket",
        "s3:ListBucketVersions"
      ],
      "Resource": ["arn:aws:s3:::amzn-s3-demo-bucket"],
      "Condition": {
        "StringEquals": {
          "aws:ResourceAccount": "123456789012"
        }
      }
    }
  ]
}

```

```

    }
  },
  {
    "Sid": "PermissionsForDecryptAnnotation",
    "Effect": "Allow",
    "Action": ["kms:Decrypt"],
    "Condition": {
      "StringLike": {
        "kms:ViaService": [
          "s3.us-east-1.amazonaws.com"
        ]
      },
      "ArnLike": {
        "kms:EncryptionContext:aws:s3:arn": [
          "arn:aws:s3:::amzn-s3-demo-bucket",
          "arn:aws:s3:::amzn-s3-demo-bucket/*"
        ]
      }
    },
    "Resource": ["arn:aws:kms:us-east-1:123456789012:key/1234abcd-12ab-34cd-56ef-1234567890ab"]
  }
]
}

```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)

- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

UpdateBucketMetadataInventoryTableConfiguration

Service: Amazon S3

Enables or disables a live inventory table for an S3 Metadata configuration on a general purpose bucket. For more information, see [Accelerating data discovery with S3 Metadata](#) in the *Amazon S3 User Guide*.

Permissions

To use this operation, you must have the following permissions. For more information, see [Setting up permissions for configuring metadata tables](#) in the *Amazon S3 User Guide*.

If you want to encrypt your inventory table with server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS), you need additional permissions in your KMS key policy. For more information, see [Setting up permissions for configuring metadata tables](#) in the *Amazon S3 User Guide*.

- `s3:UpdateBucketMetadataInventoryTableConfiguration`
- `s3tables:CreateTableBucket`
- `s3tables:CreateNamespace`
- `s3tables:GetTable`
- `s3tables:CreateTable`
- `s3tables:PutTablePolicy`
- `s3tables:PutTableEncryption`
- `kms:DescribeKey`

The following operations are related to `UpdateBucketMetadataInventoryTableConfiguration`:

- [CreateBucketMetadataConfiguration](#)
- [DeleteBucketMetadataConfiguration](#)
- [GetBucketMetadataConfiguration](#)
- [UpdateBucketMetadataJournalTableConfiguration](#)

⚠ Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /?metadataInventoryTable HTTP/1.1
Host: Bucket.s3.amazonaws.com
Content-MD5: ContentMD5
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
<InventoryTableConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <ConfigurationState>string</ConfigurationState>
  <EncryptionConfiguration>
    <KmsKeyArn>string</KmsKeyArn>
    <SseAlgorithm>string</SseAlgorithm>
  </EncryptionConfiguration>
</InventoryTableConfiguration>
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The general purpose bucket that corresponds to the metadata configuration that you want to enable or disable an inventory table for.

Required: Yes

Content-MD5

The Content-MD5 header for the inventory table configuration.

x-amz-expected-bucket-owner

The expected owner of the general purpose bucket that corresponds to the metadata table configuration that you want to enable or disable an inventory table for.

[x-amz-sdk-checksum-algorithm](#)

The checksum algorithm to use with your inventory table configuration.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 |
XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in XML format.

[InventoryTableConfiguration](#)

Root level tag for the InventoryTableConfiguration parameters.

Required: Yes

[ConfigurationState](#)

The configuration state of the inventory table, indicating whether the inventory table is enabled or disabled.

Type: String

Valid Values: ENABLED | DISABLED

Required: Yes

[EncryptionConfiguration](#)

The encryption configuration for the inventory table.

Type: [MetadataTableEncryptionConfiguration](#) data type

Required: No

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

UpdateBucketMetadataJournalTableConfiguration

Service: Amazon S3

Enables or disables journal table record expiration for an S3 Metadata configuration on a general purpose bucket. For more information, see [Accelerating data discovery with S3 Metadata](#) in the *Amazon S3 User Guide*.

Permissions

To use this operation, you must have the `s3:UpdateBucketMetadataJournalTableConfiguration` permission. For more information, see [Setting up permissions for configuring metadata tables](#) in the *Amazon S3 User Guide*.

The following operations are related to `UpdateBucketMetadataJournalTableConfiguration`:

- [CreateBucketMetadataConfiguration](#)
- [DeleteBucketMetadataConfiguration](#)
- [GetBucketMetadataConfiguration](#)
- [UpdateBucketMetadataInventoryTableConfiguration](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /?metadataJournalTable HTTP/1.1
Host: Bucket.s3.amazonaws.com
Content-MD5: ContentMD5
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-expected-bucket-owner: ExpectedBucketOwner
<?xml version="1.0" encoding="UTF-8"?>
```

```
<JournalTableConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <RecordExpiration>
    <Days>integer</Days>
    <Expiration>string</Expiration>
  </RecordExpiration>
</JournalTableConfiguration>
```

URI Request Parameters

The request uses the following URI parameters.

[Bucket](#)

The general purpose bucket that corresponds to the metadata configuration that you want to enable or disable journal table record expiration for.

Required: Yes

[Content-MD5](#)

The Content-MD5 header for the journal table configuration.

[x-amz-expected-bucket-owner](#)

The expected owner of the general purpose bucket that corresponds to the metadata table configuration that you want to enable or disable journal table record expiration for.

[x-amz-sdk-checksum-algorithm](#)

The checksum algorithm to use with your journal table configuration.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 |
XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in XML format.

[JournalTableConfiguration](#)

Root level tag for the JournalTableConfiguration parameters.

Required: Yes

[RecordExpiration](#)

The journal table record expiration settings for the journal table.

Type: [RecordExpiration](#) data type

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

UpdateObjectEncryption

Service: Amazon S3

Note

This operation is not supported for directory buckets or Amazon S3 on Outposts buckets.

Updates the server-side encryption type of an existing encrypted object in a general purpose bucket. You can use the `UpdateObjectEncryption` operation to change encrypted objects from server-side encryption with Amazon S3 managed keys (SSE-S3) to server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS), or to apply S3 Bucket Keys. You can also use the `UpdateObjectEncryption` operation to change the customer-managed KMS key used to encrypt your data so that you can comply with custom key-rotation standards.

Using the `UpdateObjectEncryption` operation, you can atomically update the server-side encryption type of an existing object in a general purpose bucket without any data movement. The `UpdateObjectEncryption` operation uses envelope encryption to re-encrypt the data key used to encrypt and decrypt your object with your newly specified server-side encryption type. In other words, when you use the `UpdateObjectEncryption` operation, your data isn't copied, archived objects in the S3 Glacier Flexible Retrieval and S3 Glacier Deep Archive storage classes aren't restored, and objects in the S3 Intelligent-Tiering storage class aren't moved between tiers. Additionally, the `UpdateObjectEncryption` operation preserves all object metadata properties, including the storage class, creation date, last modified date, ETag, and checksum properties. For more information, see [Updating server-side encryption for existing objects](#) in the *Amazon S3 User Guide*.

By default, all `UpdateObjectEncryption` requests that specify a customer-managed KMS key are restricted to KMS keys that are owned by the bucket owner's AWS account. If you're using AWS Organizations, you can request the ability to use KMS keys owned by other member accounts within your organization by contacting AWS Support.

Note

Source objects that are unencrypted, or encrypted with either dual-layer server-side encryption with AWS KMS keys (DSSE-KMS) or server-side encryption with customer-provided keys (SSE-C) aren't supported by this operation. Additionally, you cannot specify

SSE-S3 encryption as the requested new encryption type `UpdateObjectEncryption` request.

Permissions

- To use the `UpdateObjectEncryption` operation, you must have the following permissions:
 - `s3:UpdateObjectEncryption`
 - `kms:Encrypt`
 - `kms:Decrypt`
 - `kms:GenerateDataKey`
 - `kms:ReEncrypt*`
- If you're using AWS Organizations, to use this operation with customer-managed KMS keys from other AWS accounts within your organization, you must have the `organizations:DescribeAccount` permission.

Errors

- You might receive an `InvalidRequest` error for several reasons. Depending on the reason for the error, you might receive one of the following messages:
 - The `UpdateObjectEncryption` operation doesn't supported unencrypted source objects. Only source objects encrypted with SSE-S3 or SSE-KMS are supported.
 - The `UpdateObjectEncryption` operation doesn't support source objects with the encryption type DSSE-KMS or SSE-C. Only source objects encrypted with SSE-S3 or SSE-KMS are supported.
 - The `UpdateObjectEncryption` operation doesn't support updating the encryption type to DSSE-KMS or SSE-C. Modify the request to specify SSE-KMS for the updated encryption type, and then try again.
 - Requests that modify an object encryption configuration require AWS Signature Version 4. Modify the request to use AWS Signature Version 4, and then try again.
 - Requests that modify an object encryption configuration require a valid new encryption type. Valid values are SSEKMS. Modify the request to specify SSE-KMS for the updated encryption type, and then try again.

- Requests that modify an object's encryption type to SSE-KMS require an AWS KMS key Amazon Resource Name (ARN). Modify the request to specify a KMS key ARN, and then try again.
- Requests that modify an object's encryption type to SSE-KMS require a valid AWS KMS key Amazon Resource Name (ARN). Confirm that you have a correctly formatted KMS key ARN in your request, and then try again.
- The `BucketKeyEnabled` value isn't valid. Valid values are `true` or `false`. Modify the request to specify a valid value, and then try again.
- You might receive an `AccessDenied` error for several reasons. Depending on the reason for the error, you might receive one of the following messages:
 - The AWS KMS key in the request must be owned by the same account as the bucket. Modify the request to specify a KMS key from the same account, and then try again.
 - The bucket owner's account was approved to make `UpdateObjectEncryption` requests that use any AWS KMS key in their organization, but the bucket owner's account isn't part of an organization in AWS Organizations. Make sure that the bucket owner's account and the specified KMS key belong to the same organization, and then try again.
 - The specified AWS KMS key must be from the same organization in AWS Organizations as the bucket. Specify a KMS key that belongs to the same organization as the bucket, and then try again.
 - The encryption type for the specified object can't be updated because that object is protected by S3 Object Lock. If the object has a governance-mode retention period or a legal hold, you must first remove the Object Lock status on the object before you issue your `UpdateObjectEncryption` request. You can't use the `UpdateObjectEncryption` operation with objects that have an Object Lock compliance mode retention period applied to them.

Request Syntax

```
PUT /{Key+}?encryption&versionId=VersionId HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-request-payer: RequestPayer
x-amz-expected-bucket-owner: ExpectedBucketOwner
Content-MD5: ContentMD5
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
<?xml version="1.0" encoding="UTF-8"?>
<ObjectEncryption xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
```



```
<SSE-KMS>
  <BucketKeyEnabled>boolean</BucketKeyEnabled>
  <KMSKeyArn>string</KMSKeyArn>
</SSE-KMS>
</ObjectEncryption>
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the general purpose bucket that contains the specified object key name.

When you use this operation with an access point attached to a general purpose bucket, you must either provide the alias of the access point in place of the bucket name or you must specify the access point Amazon Resource Name (ARN). When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form *AccessPointName*-*AccountId*.s3-accesspoint.*Region*.amazonaws.com. When using this operation with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Referencing access points](#) in the *Amazon S3 User Guide*.

Required: Yes

Content-MD5

The MD5 hash for the request body. For requests made using the AWS Command Line Interface (CLI) or AWS SDKs, this field is calculated automatically.

Key

The key name of the object that you want to update the server-side encryption type for.

Length Constraints: Minimum length of 1.

Required: Yes

versionId

The version ID of the object that you want to update the server-side encryption type for.

x-amz-expected-bucket-owner

The account ID of the expected bucket owner. If the account ID that you provide doesn't match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

x-amz-sdk-checksum-algorithm

Indicates the algorithm used to create the checksum for the object when you use an AWS SDK. This header doesn't provide any additional functionality if you don't use the SDK. When you send this header, there must be a corresponding x-amz-checksum or x-amz-trailer header sent. Otherwise, Amazon S3 fails the request with the HTTP status code 400 Bad Request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

If you provide an individual checksum, Amazon S3 ignores any provided ChecksumAlgorithm parameter.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 | XXHASH64 | XXHASH3 | XXHASH128

Request Body

The request accepts the following data in XML format.

ObjectEncryption

Root level tag for the ObjectEncryption parameters.

Required: Yes

SSEKMS

Specifies to update the object encryption type to server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS).

Type: [SSEKMSEncryption](#) data type

Required: No

Response Syntax

```
HTTP/1.1 200
x-amz-request-charged: RequestCharged
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

Errors

AccessDenied

You might receive this error for several reasons. For details, see the description of this API operation.

HTTP Status Code: 403

InvalidRequest

A parameter or header in your request isn't valid. For details, see the description of this API operation.

HTTP Status Code: 400

NoSuchKey

The specified key does not exist.

HTTP Status Code: 404

Examples

Sample Request: Updating encryption type to SSE-KMS with S3 Bucket Keys enabled

The following example request illustrates updating the server-side encryption type of the `example.txt` object in the `amzn-s3-demo-bucket` bucket to SSE-KMS with S3 Bucket Keys enabled.

```
PUT /example.txt?encryption&versionId=VersionId HTTP/1.1
Host: amzn-s3-demo-bucket.s3.<Region>.amazonaws.com
x-amz-request-payer: <RequestPayer>
x-amz-expected-bucket-owner: <ExpectedBucketOwner>
Content-MD5: <ContentMD5>
x-amz-sdk-checksum-algorithm: <ChecksumAlgorithm>
<?xml version="1.0" encoding="UTF-8"?>
<ObjectEncryption xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <SSE-KMS>
    <BucketKeyEnabled>true</BucketKeyEnabled>
    <KMSKeyArn>arn:aws:kms:<Region>:<AccountId>:key/<KeyID></KMSKeyArn>
  </SSE-KMS>
</ObjectEncryption>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

UploadPart

Service: Amazon S3

Uploads a part in a multipart upload.

Note

In this operation, you provide new data as a part of an object in your request. However, you have an option to specify your existing Amazon S3 object as a data source for the part you are uploading. To upload a part from an existing object, you use the [UploadPartCopy](#) operation.

You must initiate a multipart upload (see [CreateMultipartUpload](#)) before you can upload any part. In response to your initiate request, Amazon S3 returns an upload ID, a unique identifier that you must include in your upload part request.

Part numbers can be any number from 1 to 10,000, inclusive. A part number uniquely identifies a part and also defines its position within the object being created. If you upload a new part using the same part number that was used with a previous part, the previously uploaded part is overwritten.

For information about maximum and minimum part sizes and other multipart upload specifications, see [Multipart upload limits](#) in the *Amazon S3 User Guide*.

Note

After you initiate multipart upload and upload one or more parts, you must either complete or abort multipart upload in order to stop getting charged for storage of the uploaded parts. Only after you either complete or abort multipart upload, Amazon S3 frees up the parts storage and stops charging you for the parts storage.

For more information on multipart uploads, go to [Multipart Upload Overview](#) in the *Amazon S3 User Guide*.

Note

Directory buckets - For directory buckets, you must make requests for this API operation to the Zonal endpoint. These endpoints support virtual-hosted-style requests in the format `https://amzn-s3-demo-bucket.s3express-zone-id.region-code.amazonaws.com/key-name` . Path-style requests are not supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

Permissions

- **General purpose bucket permissions** - To perform a multipart upload with encryption using an AWS Key Management Service key, the requester must have permission to the `kms:Decrypt` and `kms:GenerateDataKey` actions on the key. The requester must also have permissions for the `kms:GenerateDataKey` action for the `CreateMultipartUpload` API. Then, the requester needs permissions for the `kms:Decrypt` action on the `UploadPart` and `UploadPartCopy` APIs.

These permissions are required because Amazon S3 must decrypt and read data from the encrypted file parts before it completes the multipart upload. For more information about KMS permissions, see [Protecting data using server-side encryption with AWS KMS](#) in the *Amazon S3 User Guide*. For information about the permissions required to use the multipart upload API, see [Multipart upload and permissions](#) and [Multipart upload API and permissions](#) in the *Amazon S3 User Guide*.

- **Directory bucket permissions** - To grant access to this API operation on a directory bucket, we recommend that you use the [CreateSession](#) API operation for session-based authorization. Specifically, you grant the `s3express:CreateSession` permission to the directory bucket in a bucket policy or an IAM identity-based policy. Then, you make the `CreateSession` API call on the bucket to obtain a session token. With the session token in your request header, you can make API requests to this operation. After the session token expires, you make another `CreateSession` API call to generate a new session token for use. AWS CLI or SDKs create session and refresh the session token automatically to avoid service interruptions when a session expires. For more information about authorization, see [CreateSession](#).

If the object is encrypted with SSE-KMS, you must also have the `kms:GenerateDataKey` and `kms:Decrypt` permissions in IAM identity-based policies and AWS KMS key policies for the AWS KMS key.

Data integrity

General purpose bucket - To ensure that data is not corrupted traversing the network, specify the `Content-MD5` header in the upload part request. Amazon S3 checks the part data against the provided MD5 value. If they do not match, Amazon S3 returns an error. If the upload request is signed with Signature Version 4, then AWS S3 uses the `x-amz-content-sha256` header as a checksum instead of `Content-MD5`. For more information see [Authenticating Requests: Using the Authorization Header \(AWS Signature Version 4\)](#).

Note

Directory buckets - MD5 is not supported by directory buckets. You can use checksum algorithms to check object integrity.

Encryption

- **General purpose bucket** - Server-side encryption is for data encryption at rest. Amazon S3 encrypts your data as it writes it to disks in its data centers and decrypts it when you access it. You have mutually exclusive options to protect data using server-side encryption in Amazon S3, depending on how you choose to manage the encryption keys. Specifically, the encryption key options are Amazon S3 managed keys (SSE-S3), AWS KMS keys (SSE-KMS), and Customer-Provided Keys (SSE-C). Amazon S3 encrypts data with server-side encryption using Amazon S3 managed keys (SSE-S3) by default. You can optionally tell Amazon S3 to encrypt data at rest using server-side encryption with other key options. The option you use depends on whether you want to use KMS keys (SSE-KMS) or provide your own encryption key (SSE-C).

Server-side encryption is supported by the S3 Multipart Upload operations. Unless you are using a customer-provided encryption key (SSE-C), you don't need to specify the encryption parameters in each `UploadPart` request. Instead, you only need to specify the server-side encryption parameters in the initial `Initiate Multipart` request. For more information, see [CreateMultipartUpload](#).

Note

If you have server-side encryption with customer-provided keys (SSE-C) blocked for your general purpose bucket, you will get an HTTP 403 Access Denied error when you specify the SSE-C request headers while writing new data to your bucket. For more information, see [Blocking or unblocking SSE-C for a general purpose bucket](#).

If you request server-side encryption using a customer-provided encryption key (SSE-C) in your initiate multipart upload request, you must provide identical encryption information in each part upload using the following request headers.

- x-amz-server-side-encryption-customer-algorithm
- x-amz-server-side-encryption-customer-key
- x-amz-server-side-encryption-customer-key-MD5

For more information, see [Using Server-Side Encryption](#) in the *Amazon S3 User Guide*.

- **Directory buckets** - For directory buckets, there are only two supported options for server-side encryption: server-side encryption with Amazon S3 managed keys (SSE-S3) (AES256) and server-side encryption with AWS KMS keys (SSE-KMS) (aws : kms).

Special errors

- Error Code: NoSuchUpload
 - Description: The specified multipart upload does not exist. The upload ID might be invalid, or the multipart upload might have been aborted or completed.
 - HTTP Status Code: 404 Not Found
 - SOAP Fault Code Prefix: Client

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `Bucket-name.s3express-zone-id.region-code.amazonaws.com`.

The following operations are related to UploadPart:

- [CreateMultipartUpload](#)
- [CompleteMultipartUpload](#)

- [AbortMultipartUpload](#)
- [ListParts](#)
- [ListMultipartUploads](#)

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /Key?partNumber=PartNumber&uploadId=UploadId HTTP/1.1
Host: Bucket.s3.amazonaws.com
Content-Length: ContentLength
Content-MD5: ContentMD5
x-amz-sdk-checksum-algorithm: ChecksumAlgorithm
x-amz-checksum-crc32: ChecksumCRC32
x-amz-checksum-crc32c: ChecksumCRC32C
x-amz-checksum-crc64nvme: ChecksumCRC64NVME
x-amz-checksum-sha1: ChecksumSHA1
x-amz-checksum-sha256: ChecksumSHA256
x-amz-checksum-sha512: ChecksumSHA512
x-amz-checksum-md5: ChecksumMD5
x-amz-checksum-xxhash64: ChecksumXXHASH64
x-amz-checksum-xxhash3: ChecksumXXHASH3
x-amz-checksum-xxhash128: ChecksumXXHASH128
x-amz-server-side-encryption-customer-algorithm: SSECustomerAlgorithm
x-amz-server-side-encryption-customer-key: SSECustomerKey
x-amz-server-side-encryption-customer-key-MD5: SSECustomerKeyMD5
x-amz-request-payer: RequestPayer
x-amz-expected-bucket-owner: ExpectedBucketOwner
```

Body

URI Request Parameters

The request uses the following URI parameters.

Bucket

The name of the bucket to which the multipart upload was initiated.

Directory buckets - When you use this operation with a directory bucket, you must use virtual-hosted-style requests in the format `Bucket-name.s3express-zone-id.region-code.amazonaws.com`. Path-style requests are not supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must follow the format `bucket-base-name--zone-id--x-s3` (for example, `amzn-s3-demo-bucket--usw2-az1--x-s3`). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form `AccessPointName-AccountId.s3-accesspoint.Region.amazonaws.com`. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

Note

Object Lambda access points are not supported by directory buckets.

S3 on Outposts - When you use this action with S3 on Outposts, you must direct requests to the S3 on Outposts hostname. The S3 on Outposts hostname takes the form `AccessPointName-AccountId.outpostID.s3-outposts.Region.amazonaws.com`. When you use this action with S3 on Outposts, the destination bucket must be the Outposts access point ARN or the access point alias. For more information about S3 on Outposts, see [What is S3 on Outposts?](#) in the *Amazon S3 User Guide*.

Required: Yes

Content-Length

Size of the body in bytes. This parameter is useful when the size of the body cannot be determined automatically.

Content-MD5

The Base64 encoded 128-bit MD5 digest of the part data. This parameter is auto-populated when using the command from the CLI. This parameter is required if object lock parameters are specified.

Note

This functionality is not supported for directory buckets.

Key

Object key for which the multipart upload was initiated.

Length Constraints: Minimum length of 1.

Required: Yes

partNumber

Part number of part being uploaded. This is a positive integer between 1 and 10,000.

Required: Yes

uploadId

Upload ID identifying the multipart upload whose part is being uploaded.

Required: Yes

x-amz-checksum-crc32

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 32-bit CRC32 checksum of the object. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-crc32c

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 32-bit CRC32C checksum of the object. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-crc64nvme

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 64-bit CRC64NVME checksum of the part. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-md5

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 128-bit MD5 digest of the part. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-sha1

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 160-bit SHA1 digest of the object. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-sha256

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 256-bit SHA256 digest of the object. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-sha512

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 512-bit SHA512 digest of the part. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-xxhash128

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 128-bit XXHASH128 checksum of the part. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-xxhash3

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 64-bit XXHASH3 checksum of the part. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

[x-amz-checksum-xxhash64](#)

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 64-bit XXHASH64 checksum of the part. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

[x-amz-expected-bucket-owner](#)

The account ID of the expected bucket owner. If the account ID that you provide does not match the actual owner of the bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

[x-amz-request-payer](#)

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

[x-amz-sdk-checksum-algorithm](#)

Indicates the algorithm used to create the checksum for the object when you use the SDK. This header will not provide any additional functionality if you don't use the SDK. When you send this header, there must be a corresponding `x-amz-checksum` or `x-amz-trailer` header sent. Otherwise, Amazon S3 fails the request with the HTTP status code 400 Bad Request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

If you provide an individual checksum, Amazon S3 ignores any provided `ChecksumAlgorithm` parameter.

This checksum algorithm must be the same for all parts and it match the checksum value supplied in the `CreateMultipartUpload` request.

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 |
XXHASH64 | XXHASH3 | XXHASH128

[x-amz-server-side-encryption-customer-algorithm](#)

Specifies the algorithm to use when encrypting the object (for example, AES256).

Note

This functionality is not supported for directory buckets.

[x-amz-server-side-encryption-customer-key](#)

Specifies the customer-provided encryption key for Amazon S3 to use in encrypting data. This value is used to store the object and then it is discarded; Amazon S3 does not store the encryption key. The key must be appropriate for use with the algorithm specified in the `x-amz-server-side-encryption-customer-algorithm` header. This must be the same encryption key specified in the initiate multipart upload request.

Note

This functionality is not supported for directory buckets.

[x-amz-server-side-encryption-customer-key-MD5](#)

Specifies the 128-bit MD5 digest of the encryption key according to RFC 1321. Amazon S3 uses this header for a message integrity check to ensure that the encryption key was transmitted without error.

Note

This functionality is not supported for directory buckets.

Request Body

The request accepts the following binary data.

[Body](#)

Response Syntax

```
HTTP/1.1 200
x-amz-server-side-encryption: ServerSideEncryption
ETag: ETag
x-amz-checksum-crc32: ChecksumCRC32
x-amz-checksum-crc32c: ChecksumCRC32C
x-amz-checksum-crc64nvme: ChecksumCRC64NVME
x-amz-checksum-sha1: ChecksumSHA1
x-amz-checksum-sha256: ChecksumSHA256
x-amz-checksum-sha512: ChecksumSHA512
x-amz-checksum-md5: ChecksumMD5
x-amz-checksum-xxhash64: ChecksumXXHASH64
x-amz-checksum-xxhash3: ChecksumXXHASH3
x-amz-checksum-xxhash128: ChecksumXXHASH128
x-amz-server-side-encryption-customer-algorithm: SSECustomerAlgorithm
x-amz-server-side-encryption-customer-key-MD5: SSECustomerKeyMD5
x-amz-server-side-encryption-aws-kms-key-id: SSEKMSKeyId
x-amz-server-side-encryption-bucket-key-enabled: BucketKeyEnabled
x-amz-request-charged: RequestCharged
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

ETag

Entity tag for the uploaded object.

x-amz-checksum-crc32

The Base64 encoded, 32-bit CRC32 checksum of the part. This will only be present if the checksum was provided in the request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-crc32c

The Base64 encoded, 32-bit CRC32C checksum of the part. This will only be present if the checksum was provided in the request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-crc64nvme

The Base64 encoded, 64-bit CRC64NVME checksum of the part. This will only be present if the checksum was provided in the request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-md5

The Base64 encoded, 128-bit MD5 checksum of the part. This will only be present if the checksum was provided in the request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-sha1

The Base64 encoded, 160-bit SHA1 checksum of the part. This will only be present if the checksum was provided in the request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-sha256

The Base64 encoded, 256-bit SHA256 checksum of the part. This will only be present if the checksum was provided in the request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-sha512

The Base64 encoded, 512-bit SHA512 checksum of the part. This will only be present if the checksum was provided in the request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-xxhash128

The Base64 encoded, 128-bit XXHASH128 checksum of the part. This will only be present if the checksum was provided in the request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-xxhash3

The Base64 encoded, 64-bit XXHASH3 checksum of the part. This will only be present if the checksum was provided in the request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-checksum-xxhash64

The Base64 encoded, 64-bit XXHASH64 checksum of the part. This will only be present if the checksum was provided in the request. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

x-amz-server-side-encryption

The server-side encryption algorithm used when you store this object in Amazon S3 or Amazon FSx.

Note

When accessing data stored in Amazon FSx file systems using S3 access points, the only valid server side encryption option is `aws:fsx`.

Valid Values: AES256 | `aws:fsx` | `aws:backup` | `aws:kms` | `aws:kms:dsse`

x-amz-server-side-encryption-aws-kms-key-id

If present, indicates the ID of the KMS key that was used for object encryption.

x-amz-server-side-encryption-bucket-key-enabled

Indicates whether the multipart upload uses an S3 Bucket Key for server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS).

x-amz-server-side-encryption-customer-algorithm

If server-side encryption with a customer-provided encryption key was requested, the response will include this header to confirm the encryption algorithm that's used.

Note

This functionality is not supported for directory buckets.

x-amz-server-side-encryption-customer-key-MD5

If server-side encryption with a customer-provided encryption key was requested, the response will include this header to provide the round-trip message integrity verification of the customer-provided encryption key.

Note

This functionality is not supported for directory buckets.

Examples

Sample Request for general purpose buckets

The following PUT request uploads a part (part number 1) in a multipart upload. The request includes the upload ID that you get in response to your Initiate Multipart Upload request.

```
PUT /my-movie.m2ts?
partNumber=1&uploadId=VCVsb2FkIElEIGZvciBlbZZpbmcncyBteS1tb3ZpZS5tMnRzIHVwbG9hZR
HTTP/1.1
Host: example-bucket.s3.<Region>.amazonaws.com
Date: Mon, 1 Nov 2010 20:34:56 GMT
Content-Length: 10485760
Content-MD5: pUNXr/BjKK5G2UKvaRRr0A==
Authorization: authorization string

***part data omitted***
```

Sample Response for general purpose buckets

The response includes the ETag header. You need to retain this value for use when you send the Complete Multipart Upload request.

```
HTTP/1.1 200 OK
x-amz-id-2: Vvag1LuByRx9e6j50nimru9p04ZVKnJ2Qz7/C1NPcfTWAtRPfTa0Fg==
x-amz-request-id: 656c76696e6727732072657175657374
Date: Mon, 1 Nov 2010 20:34:56 GMT
ETag: "b54357faf0632cce46e942fa68356b38"
Content-Length: 0
Connection: keep-alive
Server: AmazonS3
```

Example for general purpose buckets: Upload a part with an encryption key in the request for server-side encryption

If you initiated a multipart upload with a request to save an object using server-side encryption with a customer-provided encryption key, each part upload must also include the same set of encryption-specific headers as shown in the following example request.

```
PUT /example-object?
partNumber=1&uploadId=EXAMPLEJZ6e0YupT2h66iePQCc9IEbYbDUy4RTpMeoSMLPRp8Z5o1u8feSRonpvnWsKKG35tI
HTTP/1.1
Host: example-bucket.s3.<Region>.amazonaws.com
Authorization: authorization string
Date: Wed, 28 May 2014 19:40:11 +0000
x-amz-server-side-encryption-customer-key: g0lCfA3Dv40jZz5SQJ1ZukLRFqtI5WorC/8SEEXAMPLE

x-amz-server-side-encryption-customer-key-MD5: ZjQrne1X/iTcskbY2example
x-amz-server-side-encryption-customer-algorithm: AES256
```

Example for general purpose buckets

In the response, Amazon S3 returns encryption-specific headers providing the encryption algorithm used and MD5 digest of the encryption key you provided in the request.

```
HTTP/1.1 100 Continue    HTTP/1.1 200 OK
x-amz-id-2: Zn8bf8aEFQ+kBnGPBc/JaAf9SoWM68QDPS9+SyFwkIZ0HUG2BiRLZi5oXw4c0CEt
x-amz-request-id: 5A37448A37622243
Date: Wed, 28 May 2014 19:40:12 GMT
ETag: "7e10e7d25dc4581d89b9285be5f384fd"
x-amz-server-side-encryption-customer-algorithm: AES256
x-amz-server-side-encryption-customer-key-MD5: ZjQrne1X/iTcskbY2example
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

UploadPartCopy

Service: Amazon S3

Uploads a part by copying data from an existing object as data source. To specify the data source, you add the request header `x-amz-copy-source` in your request. To specify a byte range, you add the request header `x-amz-copy-source-range` in your request.

For information about maximum and minimum part sizes and other multipart upload specifications, see [Multipart upload limits](#) in the *Amazon S3 User Guide*.

Note

Instead of copying data from an existing object as part data, you might use the [UploadPart](#) action to upload new data as a part of an object in your request.

You must initiate a multipart upload before you can upload any part. In response to your initiate request, Amazon S3 returns the upload ID, a unique identifier that you must include in your upload part request.

For conceptual information about multipart uploads, see [Uploading Objects Using Multipart Upload](#) in the *Amazon S3 User Guide*. For information about copying objects using a single atomic action vs. a multipart upload, see [Operations on Objects](#) in the *Amazon S3 User Guide*.

Note

Directory buckets - For directory buckets, you must make requests for this API operation to the Zonal endpoint. These endpoints support virtual-hosted-style requests in the format `https://amzn-s3-demo-bucket.s3express-zone-id.region-code.amazonaws.com/key-name`. Path-style requests are not supported. For more information about endpoints in Availability Zones, see [Regional and Zonal endpoints for directory buckets in Availability Zones](#) in the *Amazon S3 User Guide*. For more information about endpoints in Local Zones, see [Concepts for directory buckets in Local Zones](#) in the *Amazon S3 User Guide*.

Authentication and authorization

All UploadPartCopy requests must be authenticated and signed by using IAM credentials (access key ID and secret access key for the IAM identities). All headers with the x-amz- prefix, including x-amz-copy-source, must be signed. For more information, see [REST Authentication](#).

Directory buckets - You must use IAM credentials to authenticate and authorize your access to the UploadPartCopy API operation, instead of using the temporary security credentials through the CreateSession API operation.

AWS CLI or SDKs handles authentication and authorization on your behalf.

Permissions

You must have READ access to the source object and WRITE access to the destination bucket.

- **General purpose bucket permissions** - You must have the permissions in a policy based on the bucket types of your source bucket and destination bucket in an UploadPartCopy operation.
 - If the source object is in a general purpose bucket, you must have the **s3:GetObject** permission to read the source object that is being copied.
 - If the destination bucket is a general purpose bucket, you must have the **s3:PutObject** permission to write the object copy to the destination bucket.
 - To perform a multipart upload with encryption using an AWS Key Management Service key, the requester must have permission to the kms:Decrypt and kms:GenerateDataKey actions on the key. The requester must also have permissions for the kms:GenerateDataKey action for the CreateMultipartUpload API. Then, the requester needs permissions for the kms:Decrypt action on the UploadPart and UploadPartCopy APIs. These permissions are required because Amazon S3 must decrypt and read data from the encrypted file parts before it completes the multipart upload. For more information about KMS permissions, see [Protecting data using server-side encryption with AWS KMS](#) in the *Amazon S3 User Guide*. For information about the permissions required to use the multipart upload API, see [Multipart upload and permissions](#) and [Multipart upload API and permissions](#) in the *Amazon S3 User Guide*.
- **Directory bucket permissions** - You must have permissions in a bucket policy or an IAM identity-based policy based on the source and destination bucket types in an UploadPartCopy operation.

- If the source object that you want to copy is in a directory bucket, you must have the **s3express:CreateSession** permission in the Action element of a policy to read the object. If no session mode is specified, the session will be created with the maximum allowable privilege, attempting ReadWrite first, then ReadOnly if ReadWrite is not permitted. If you want to explicitly restrict the access to be read-only, you can set the **s3express:SessionMode** condition key to ReadOnly on the copy source bucket.
- If the copy destination is a directory bucket, you must have the **s3express:CreateSession** permission in the Action element of a policy to write the object to the destination. The **s3express:SessionMode** condition key cannot be set to ReadOnly on the copy destination.

If the object is encrypted with SSE-KMS, you must also have the **kms:GenerateDataKey** and **kms:Decrypt** permissions in IAM identity-based policies and AWS KMS key policies for the AWS KMS key.

For example policies, see [Example bucket policies for S3 Express One Zone](#) and [AWS Identity and Access Management \(IAM\) identity-based policies for S3 Express One Zone](#) in the *Amazon S3 User Guide*.

Encryption

- **General purpose buckets** - For information about using server-side encryption with customer-provided encryption keys with the UploadPartCopy operation, see [CopyObject](#) and [UploadPart](#).

Note

If you have server-side encryption with customer-provided keys (SSE-C) blocked for your general purpose bucket, you will get an HTTP 403 Access Denied error when you specify the SSE-C request headers while writing new data to your bucket. For more information, see [Blocking or unblocking SSE-C for a general purpose bucket](#).

- **Directory buckets** - For directory buckets, there are only two supported options for server-side encryption: server-side encryption with Amazon S3 managed keys (SSE-S3) (AES256) and server-side encryption with AWS KMS keys (SSE-KMS) (**aws:kms**). For more information, see [Protecting data with server-side encryption](#) in the *Amazon S3 User Guide*.

Note

For directory buckets, when you perform a `CreateMultipartUpload` operation and an `UploadPartCopy` operation, the request headers you provide in the `CreateMultipartUpload` request must match the default encryption configuration of the destination bucket.

S3 Bucket Keys aren't supported, when you copy SSE-KMS encrypted objects from general purpose buckets to directory buckets, from directory buckets to general purpose buckets, or between directory buckets, through [UploadPartCopy](#). In this case, Amazon S3 makes a call to AWS KMS every time a copy request is made for a KMS-encrypted object.

Special errors

- Error Code: `NoSuchUpload`
 - Description: The specified multipart upload does not exist. The upload ID might be invalid, or the multipart upload might have been aborted or completed.
 - HTTP Status Code: 404 Not Found
- Error Code: `InvalidRequest`
 - Description: The specified copy source is not supported as a byte-range copy source.
 - HTTP Status Code: 400 Bad Request

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `Bucket-name.s3express-zone-id.region-code.amazonaws.com`.

The following operations are related to `UploadPartCopy`:

- [CreateMultipartUpload](#)
- [UploadPart](#)
- [CompleteMultipartUpload](#)
- [AbortMultipartUpload](#)
- [ListParts](#)
- [ListMultipartUploads](#)

⚠ Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
PUT /Key?partNumber=PartNumber&uploadId=UploadId HTTP/1.1
Host: Bucket.s3.amazonaws.com
x-amz-copy-source: CopySource
x-amz-copy-source-if-match: CopySourceIfMatch
x-amz-copy-source-if-modified-since: CopySourceIfModifiedSince
x-amz-copy-source-if-none-match: CopySourceIfNoneMatch
x-amz-copy-source-if-unmodified-since: CopySourceIfUnmodifiedSince
x-amz-copy-source-range: CopySourceRange
x-amz-server-side-encryption-customer-algorithm: SSECustomerAlgorithm
x-amz-server-side-encryption-customer-key: SSECustomerKey
x-amz-server-side-encryption-customer-key-MD5: SSECustomerKeyMD5
x-amz-copy-source-server-side-encryption-customer-
algorithm: CopySourceSSECustomerAlgorithm
x-amz-copy-source-server-side-encryption-customer-key: CopySourceSSECustomerKey
x-amz-copy-source-server-side-encryption-customer-key-MD5: CopySourceSSECustomerKeyMD5
x-amz-request-payer: RequestPayer
x-amz-expected-bucket-owner: ExpectedBucketOwner
x-amz-source-expected-bucket-owner: ExpectedSourceBucketOwner
```

URI Request Parameters

The request uses the following URI parameters.

Bucket

The bucket name.

Directory buckets - When you use this operation with a directory bucket, you must use virtual-hosted-style requests in the format `Bucket-name.s3express-zone-id.region-code.amazonaws.com`. Path-style requests are not supported. Directory bucket names must be unique in the chosen Zone (Availability Zone or Local Zone). Bucket names must follow the format `bucket-base-name--zone-id--x-s3` (for example, `amzn-s3-demo-`

`bucket--usw2-az1--x-s3`). For information about bucket naming restrictions, see [Directory bucket naming rules](#) in the *Amazon S3 User Guide*.

 Note

Copying objects across different AWS Regions isn't supported when the source or destination bucket is in AWS Local Zones. The source and destination buckets must have the same parent AWS Region. Otherwise, you get an HTTP 400 Bad Request error with the error code `InvalidRequest`.

Access points - When you use this action with an access point for general purpose buckets, you must provide the alias of the access point in place of the bucket name or specify the access point ARN. When you use this action with an access point for directory buckets, you must provide the access point name in place of the bucket name. When using the access point ARN, you must direct requests to the access point hostname. The access point hostname takes the form `AccessPointName-AccountId.s3-accesspoint.Region.amazonaws.com`. When using this action with an access point through the AWS SDKs, you provide the access point ARN in place of the bucket name. For more information about access point ARNs, see [Using access points](#) in the *Amazon S3 User Guide*.

 Note

Object Lambda access points are not supported by directory buckets.

S3 on Outposts - When you use this action with S3 on Outposts, you must direct requests to the S3 on Outposts hostname. The S3 on Outposts hostname takes the form `AccessPointName-AccountId.outpostID.s3-outposts.Region.amazonaws.com`. When you use this action with S3 on Outposts, the destination bucket must be the Outposts access point ARN or the access point alias. For more information about S3 on Outposts, see [What is S3 on Outposts?](#) in the *Amazon S3 User Guide*.

Required: Yes

Key

Object key for which the multipart upload was initiated.

Length Constraints: Minimum length of 1.

Required: Yes

partNumber

Part number of part being copied. This is a positive integer between 1 and 10,000.

Required: Yes

uploadId

Upload ID identifying the multipart upload whose part is being copied.

Required: Yes

x-amz-copy-source

Specifies the source object for the copy operation. You specify the value in one of two formats, depending on whether you want to access the source object through an [access point](#):

- For objects not accessed through an access point, specify the name of the source bucket and key of the source object, separated by a slash (/). For example, to copy the object `reports/january.pdf` from the bucket `awsexamplebucket`, use `awsexamplebucket/reports/january.pdf`. The value must be URL-encoded.
- For objects accessed through access points, specify the Amazon Resource Name (ARN) of the object as accessed through the access point, in the format `arn:aws:s3:<Region>:<account-id>:accesspoint/<access-point-name>/object/<key>`. For example, to copy the object `reports/january.pdf` through access point `my-access-point` owned by account `123456789012` in Region `us-west-2`, use the URL encoding of `arn:aws:s3:us-west-2:123456789012:accesspoint/my-access-point/object/reports/january.pdf`. The value must be URL encoded.

Note

- Amazon S3 supports copy operations using Access points only when the source and destination buckets are in the same AWS Region.
- Access points are not supported by directory buckets.

Alternatively, for objects accessed through Amazon S3 on Outposts, specify the ARN of the object as accessed in the format `arn:aws:s3-outposts:<Region>:<account-id>:outpost/<outpost-id>/object/<key>`. For example, to copy the object `reports/january.pdf` through outpost `my-outpost` owned by account `123456789012`

in Region us-west-2, use the URL encoding of `arn:aws:s3-outposts:us-west-2:123456789012:outpost/my-outpost/object/reports/january.pdf`. The value must be URL-encoded.

If your bucket has versioning enabled, you could have multiple versions of the same object. By default, `x-amz-copy-source` identifies the current version of the source object to copy. To copy a specific version of the source object to copy, append `?versionId=<version-id>` to the `x-amz-copy-source` request header (for example, `x-amz-copy-source: /awsexamplebucket/reports/january.pdf?versionId=QUpfndhfd8438MNFDN93jdnJFkdmqnh893`).

If the current version is a delete marker and you don't specify a `versionId` in the `x-amz-copy-source` request header, Amazon S3 returns a `404 Not Found` error, because the object does not exist. If you specify `versionId` in the `x-amz-copy-source` and the `versionId` is a delete marker, Amazon S3 returns an `HTTP 400 Bad Request` error, because you are not allowed to specify a delete marker as a version for the `x-amz-copy-source`.

 Note

Directory buckets - S3 Versioning isn't enabled and supported for directory buckets.

Pattern: `\/?\.\+\/\.\+`

Required: Yes

[x-amz-copy-source-if-match](#)

Copies the object if its entity tag (ETag) matches the specified tag.

If both of the `x-amz-copy-source-if-match` and `x-amz-copy-source-if-unmodified-since` headers are present in the request as follows:

`x-amz-copy-source-if-match` condition evaluates to true, and;

`x-amz-copy-source-if-unmodified-since` condition evaluates to false;

Amazon S3 returns `200 OK` and copies the data.

[x-amz-copy-source-if-modified-since](#)

Copies the object if it has been modified since the specified time.

If both of the `x-amz-copy-source-if-none-match` and `x-amz-copy-source-if-modified-since` headers are present in the request as follows:

`x-amz-copy-source-if-none-match` condition evaluates to false, and;

`x-amz-copy-source-if-modified-since` condition evaluates to true;

Amazon S3 returns 412 `Precondition Failed` response code.

`x-amz-copy-source-if-none-match`

Copies the object if its entity tag (ETag) is different than the specified ETag.

If both of the `x-amz-copy-source-if-none-match` and `x-amz-copy-source-if-modified-since` headers are present in the request as follows:

`x-amz-copy-source-if-none-match` condition evaluates to false, and;

`x-amz-copy-source-if-modified-since` condition evaluates to true;

Amazon S3 returns 412 `Precondition Failed` response code.

`x-amz-copy-source-if-unmodified-since`

Copies the object if it hasn't been modified since the specified time.

If both of the `x-amz-copy-source-if-match` and `x-amz-copy-source-if-unmodified-since` headers are present in the request as follows:

`x-amz-copy-source-if-match` condition evaluates to true, and;

`x-amz-copy-source-if-unmodified-since` condition evaluates to false;

Amazon S3 returns 200 `OK` and copies the data.

`x-amz-copy-source-range`

The range of bytes to copy from the source object. The range value must use the form `bytes=first-last`, where the first and last are the zero-based byte offsets to copy. For example, `bytes=0-9` indicates that you want to copy the first 10 bytes of the source. You can copy a range only if the source object is greater than 5 MB.

`x-amz-copy-source-server-side-encryption-customer-algorithm`

Specifies the algorithm to use when decrypting the source object (for example, AES256).

Note

This functionality is not supported when the source object is in a directory bucket.

x-amz-copy-source-server-side-encryption-customer-key

Specifies the customer-provided encryption key for Amazon S3 to use to decrypt the source object. The encryption key provided in this header must be one that was used when the source object was created.

Note

This functionality is not supported when the source object is in a directory bucket.

x-amz-copy-source-server-side-encryption-customer-key-MD5

Specifies the 128-bit MD5 digest of the encryption key according to RFC 1321. Amazon S3 uses this header for a message integrity check to ensure that the encryption key was transmitted without error.

Note

This functionality is not supported when the source object is in a directory bucket.

x-amz-expected-bucket-owner

The account ID of the expected destination bucket owner. If the account ID that you provide does not match the actual owner of the destination bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

x-amz-request-payer

Confirms that the requester knows that they will be charged for the request. Bucket owners need not specify this parameter in their requests. If either the source or destination S3 bucket has Requester Pays enabled, the requester will pay for the corresponding charges. For information about downloading objects from Requester Pays buckets, see [Downloading Objects in Requester Pays Buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

x-amz-server-side-encryption-customer-algorithm

Specifies the algorithm to use when encrypting the object (for example, AES256).

Note

This functionality is not supported when the destination bucket is a directory bucket.

x-amz-server-side-encryption-customer-key

Specifies the customer-provided encryption key for Amazon S3 to use in encrypting data. This value is used to store the object and then it is discarded; Amazon S3 does not store the encryption key. The key must be appropriate for use with the algorithm specified in the `x-amz-server-side-encryption-customer-algorithm` header. This must be the same encryption key specified in the initiate multipart upload request.

Note

This functionality is not supported when the destination bucket is a directory bucket.

x-amz-server-side-encryption-customer-key-MD5

Specifies the 128-bit MD5 digest of the encryption key according to RFC 1321. Amazon S3 uses this header for a message integrity check to ensure that the encryption key was transmitted without error.

Note

This functionality is not supported when the destination bucket is a directory bucket.

x-amz-source-expected-bucket-owner

The account ID of the expected source bucket owner. If the account ID that you provide does not match the actual owner of the source bucket, the request fails with the HTTP status code 403 Forbidden (access denied).

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
x-amz-copy-source-version-id: CopySourceVersionId
x-amz-server-side-encryption: ServerSideEncryption
x-amz-server-side-encryption-customer-algorithm: SSECustomerAlgorithm
x-amz-server-side-encryption-customer-key-MD5: SSECustomerKeyMD5
x-amz-server-side-encryption-aws-kms-key-id: SSEKMSKeyId
x-amz-server-side-encryption-bucket-key-enabled: BucketKeyEnabled
x-amz-request-charged: RequestCharged
<?xml version="1.0" encoding="UTF-8"?>
<CopyPartResult>
  <ETag>string</ETag>
  <LastModified>timestamp</LastModified>
  <ChecksumCRC32>string</ChecksumCRC32>
  <ChecksumCRC32C>string</ChecksumCRC32C>
  <ChecksumCRC64NVME>string</ChecksumCRC64NVME>
  <ChecksumSHA1>string</ChecksumSHA1>
  <ChecksumSHA256>string</ChecksumSHA256>
  <ChecksumSHA512>string</ChecksumSHA512>
  <ChecksumMD5>string</ChecksumMD5>
  <ChecksumXXHASH64>string</ChecksumXXHASH64>
  <ChecksumXXHASH3>string</ChecksumXXHASH3>
  <ChecksumXXHASH128>string</ChecksumXXHASH128>
</CopyPartResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

x-amz-copy-source-version-id

The version of the source object that was copied, if you have enabled versioning on the source bucket.

Note

This functionality is not supported when the source object is in a directory bucket.

x-amz-request-charged

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: requester

x-amz-server-side-encryption

The server-side encryption algorithm used when you store this object in Amazon S3 or Amazon FSx.

Note

When accessing data stored in Amazon FSx file systems using S3 access points, the only valid server side encryption option is `aws:fsx`.

Valid Values: AES256 | `aws:fsx` | `aws:backup` | `aws:kms` | `aws:kms:dsse`

x-amz-server-side-encryption-aws-kms-key-id

If present, indicates the ID of the KMS key that was used for object encryption.

[x-amz-server-side-encryption-bucket-key-enabled](#)

Indicates whether the multipart upload uses an S3 Bucket Key for server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS).

[x-amz-server-side-encryption-customer-algorithm](#)

If server-side encryption with a customer-provided encryption key was requested, the response will include this header to confirm the encryption algorithm that's used.

Note

This functionality is not supported for directory buckets.

[x-amz-server-side-encryption-customer-key-MD5](#)

If server-side encryption with a customer-provided encryption key was requested, the response will include this header to provide the round-trip message integrity verification of the customer-provided encryption key.

Note

This functionality is not supported for directory buckets.

The following data is returned in XML format by the service.

[CopyPartResult](#)

Root level tag for the CopyPartResult parameters.

Required: Yes

[ChecksumCRC32](#)

The Base64 encoded, 32-bit CRC32 checksum of the part. This checksum is present if the multipart upload request was created with the CRC32 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

ChecksumCRC32C

The Base64 encoded, 32-bit CRC32C checksum of the part. This checksum is present if the multipart upload request was created with the CRC32C checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

ChecksumCRC64NVME

The Base64 encoded, 64-bit CRC64NVME checksum of the part. This checksum is present if the multipart upload request was created with the CRC64NVME checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

ChecksumMD5

The Base64 encoded, 128-bit MD5 digest of the part. This checksum is present if the multipart upload request was created with the MD5 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

ChecksumSHA1

The Base64 encoded, 160-bit SHA1 digest of the part. This checksum is present if the multipart upload request was created with the SHA1 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

ChecksumSHA256

The Base64 encoded, 256-bit SHA256 digest of the part. This checksum is present if the multipart upload request was created with the SHA256 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

ChecksumSHA512

The Base64 encoded, 512-bit SHA512 digest of the part. This checksum is present if the multipart upload request was created with the SHA512 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

ChecksumXXHASH128

The Base64 encoded, 128-bit XXHASH128 checksum of the part. This checksum is present if the multipart upload request was created with the XXHASH128 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

ChecksumXXHASH3

The Base64 encoded, 64-bit XXHASH3 checksum of the part. This checksum is present if the multipart upload request was created with the XXHASH3 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

ChecksumXXHASH64

The Base64 encoded, 64-bit XXHASH64 checksum of the part. This checksum is present if the multipart upload request was created with the XXHASH64 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

ETag

Entity tag of the object.

Type: String

LastModified

Date and time at which the object was uploaded.

Type: Timestamp

Examples

Sample Request for general purpose buckets

The following PUT request uploads a part (part number 2) in a multipart upload. The request specifies a byte range from an existing object as the source of this upload. The request includes the upload ID that you get in response to your Initiate Multipart Upload request.

```
PUT /newobject?
partNumber=2&uploadId=VCVsb2FkIElEIGZvciBlbZZpbmcncyBteS1tb3ZpZS5tMnRzIHVwbG9hZR
HTTP/1.1
Host: target-bucket.s3.<Region>.amazonaws.com
Date: Mon, 11 Apr 2011 20:34:56 GMT
x-amz-copy-source: /source-bucket/sourceobject
x-amz-copy-source-range:bytes=500-6291456
Authorization: authorization string
```

Sample Response for general purpose buckets

The response includes the ETag value. You need to retain this value to use when you send the Complete Multipart Upload request.

```
HTTP/1.1 200 OK
x-amz-id-2: Vvag1LuByRx9e6j50nimru9p04ZVKnJ2Qz7/C1NPcfTWAtRPfTa0Fg==
x-amz-request-id: 656c76696e6727732072657175657374
Date: Mon, 11 Apr 2011 20:34:56 GMT
Server: AmazonS3

<CopyPartResult>
  <LastModified>2011-04-11T20:34:56.000Z</LastModified>
  <ETag>"9b2cf535f27731c974343645a3985328"</ETag>
</CopyPartResult>
```

Sample Request for general purpose buckets

The following PUT request uploads a part (part number 2) in a multipart upload. The request does not specify the optional byte range header, but requests the entire source object copy as part 2. The request includes the upload ID that you got in response to your Initiate Multipart Upload request.

```
PUT /newobject?
partNumber=2&uploadId=VCVsb2FkIElEIGZvciBlbZZpbmcncyBteS1tb3ZpZS5tMnRzIHVwbG9hZR
HTTP/1.1
Host: target-bucket.s3.<Region>.amazonaws.com
```

```
Date: Mon, 11 Apr 2011 20:34:56 GMT
x-amz-copy-source: /source-bucket/sourceobject?versionId=3/L4kqtJlcpXroDTDmJ
+rmSpXd3dIbrHY+MTRCxf3vjVBH40Nr8X8gdRQBpUMLUo
Authorization: authorization string
```

Sample Response for general purpose buckets

The response includes the ETag value. You need to retain this value to use when you send the Complete Multipart Upload request.

```
HTTP/1.1 200 OK
x-amz-id-2: Vvag1LuByRx9e6j50nimru9p04ZVKnJ2Qz7/C1NPcfTWAtRPfTa0Fg==
x-amz-request-id: 656c76696e6727732072657175657374
x-amz-copy-source-version-id: 3/L4kqtJlcpXroDTDmJ+rmSpXd3dIbrHY
+MTRCxf3vjVBH40Nr8X8gdRQBpUMLUo
Date: Mon, 11 Apr 2011 20:34:56 GMT
Server: AmazonS3

<CopyPartResult>
  <LastModified>2011-04-11T20:34:56.000Z</LastModified>
  <ETag>"9b2cf535f27731c974343645a3985328"</ETag>
</CopyPartResult>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)

- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

WriteGetObjectResponse

Service: Amazon S3

Note

This operation is not supported for directory buckets.

Passes transformed objects to a `GetObject` operation when using Object Lambda access points. For information about Object Lambda access points, see [Transforming objects with Object Lambda access points](#) in the *Amazon S3 User Guide*.

This operation supports metadata that can be returned by [GetObject](#), in addition to `RequestRoute`, `RequestToken`, `StatusCode`, `ErrorCode`, and `ErrorMessage`. The `GetObject` response metadata is supported so that the `WriteGetObjectResponse` caller, typically an AWS Lambda function, can provide the same metadata when it internally invokes `GetObject`. When `WriteGetObjectResponse` is called by a customer-owned Lambda function, the metadata returned to the end user `GetObject` call might differ from what Amazon S3 would normally return.

You can include any number of metadata headers. When including a metadata header, it should be prefaced with `x-amz-meta`. For example, `x-amz-meta-my-custom-header: MyCustomValue`. The primary use case for this is to forward `GetObject` metadata.

AWS provides some prebuilt Lambda functions that you can use with S3 Object Lambda to detect and redact personally identifiable information (PII) and decompress S3 objects. These Lambda functions are available in the AWS Serverless Application Repository, and can be selected through the AWS Management Console when you create your Object Lambda access point.

Example 1: PII Access Control - This Lambda function uses Amazon Comprehend, a natural language processing (NLP) service using machine learning to find insights and relationships in text. It automatically detects personally identifiable information (PII) such as names, addresses, dates, credit card numbers, and social security numbers from documents in your Amazon S3 bucket.

Example 2: PII Redaction - This Lambda function uses Amazon Comprehend, a natural language processing (NLP) service using machine learning to find insights and relationships in text. It automatically redacts personally identifiable information (PII) such as names, addresses, dates, credit card numbers, and social security numbers from documents in your Amazon S3 bucket.

Example 3: Decompression - The Lambda function `S3ObjectLambdaDecompression`, is equipped to decompress objects stored in S3 in one of six compressed file formats including bzip2, gzip, snappy, zlib, zstandard and ZIP.

For information on how to view and use these functions, see [Using AWS built Lambda functions](#) in the *Amazon S3 User Guide*.

Important

You must URL encode any signed header values that contain spaces. For example, if your header value is `my file.txt`, containing two spaces after `my`, you must URL encode this value to `my%20%20file.txt`.

Request Syntax

```
POST /WriteGetObjectResponse HTTP/1.1
Host: s3.amazonaws.com
x-amz-request-route: RequestRoute
x-amz-request-token: RequestToken
x-amz-fwd-status: StatusCode
x-amz-fwd-error-code: ErrorCode
x-amz-fwd-error-message: ErrorMessage
x-amz-fwd-header-accept-ranges: AcceptRanges
x-amz-fwd-header-Cache-Control: CacheControl
x-amz-fwd-header-Content-Disposition: ContentDisposition
x-amz-fwd-header-Content-Encoding: ContentEncoding
x-amz-fwd-header-Content-Language: ContentLanguage
Content-Length: ContentLength
x-amz-fwd-header-Content-Range: ContentRange
x-amz-fwd-header-Content-Type: ContentType
x-amz-fwd-header-x-amz-checksum-crc32: ChecksumCRC32
x-amz-fwd-header-x-amz-checksum-crc32c: ChecksumCRC32C
x-amz-fwd-header-x-amz-checksum-crc64nvme: ChecksumCRC64NVME
x-amz-fwd-header-x-amz-checksum-sha1: ChecksumSHA1
x-amz-fwd-header-x-amz-checksum-sha256: ChecksumSHA256
x-amz-fwd-header-x-amz-checksum-sha512: ChecksumSHA512
x-amz-fwd-header-x-amz-checksum-md5: ChecksumMD5
x-amz-fwd-header-x-amz-checksum-xxhash64: ChecksumXXHASH64
x-amz-fwd-header-x-amz-checksum-xxhash3: ChecksumXXHASH3
x-amz-fwd-header-x-amz-checksum-xxhash128: ChecksumXXHASH128
x-amz-fwd-header-x-amz-delete-marker: DeleteMarker
```

```

x-amz-fwd-header-ETag: ETag
x-amz-fwd-header-Expires: Expires
x-amz-fwd-header-x-amz-expiration: Expiration
x-amz-fwd-header-Last-Modified: LastModified
x-amz-fwd-header-x-amz-missing-meta: MissingMeta
x-amz-fwd-header-x-amz-object-lock-mode: ObjectLockMode
x-amz-fwd-header-x-amz-object-lock-legal-hold: ObjectLockLegalHoldStatus
x-amz-fwd-header-x-amz-object-lock-retain-until-date: ObjectLockRetainUntilDate
x-amz-fwd-header-x-amz-mp-parts-count: PartsCount
x-amz-fwd-header-x-amz-replication-status: ReplicationStatus
x-amz-fwd-header-x-amz-request-charged: RequestCharged
x-amz-fwd-header-x-amz-restore: Restore
x-amz-fwd-header-x-amz-server-side-encryption: ServerSideEncryption
x-amz-fwd-header-x-amz-server-side-encryption-customer-algorithm: SSECustomerAlgorithm
x-amz-fwd-header-x-amz-server-side-encryption-aws-kms-key-id: SSEKMSKeyId
x-amz-fwd-header-x-amz-server-side-encryption-customer-key-MD5: SSECustomerKeyMD5
x-amz-fwd-header-x-amz-storage-class: StorageClass
x-amz-fwd-header-x-amz-tagging-count: TagCount
x-amz-fwd-header-x-amz-version-id: VersionId
x-amz-fwd-header-x-amz-server-side-encryption-bucket-key-enabled: BucketKeyEnabled

```

Body

URI Request Parameters

The request uses the following URI parameters.

Content-Length

The size of the content body in bytes.

x-amz-fwd-error-code

A string that uniquely identifies an error condition. Returned in the <Code> tag of the error XML response for a corresponding GetObject call. Cannot be used with a successful StatusCode header or when the transformed object is provided in the body. All error codes from S3 are sentence-cased. The regular expression (regex) value is "[A-Z][a-zA-Z]+\$".

x-amz-fwd-error-message

Contains a generic description of the error condition. Returned in the <Message> tag of the error XML response for a corresponding GetObject call. Cannot be used with a successful StatusCode header or when the transformed object is provided in body.

x-amz-fwd-header-accept-ranges

Indicates that a range of bytes was specified.

x-amz-fwd-header-Cache-Control

Specifies caching behavior along the request/reply chain.

x-amz-fwd-header-Content-Disposition

Specifies presentational information for the object.

x-amz-fwd-header-Content-Encoding

Specifies what content encodings have been applied to the object and thus what decoding mechanisms must be applied to obtain the media-type referenced by the Content-Type header field.

x-amz-fwd-header-Content-Language

The language the content is in.

x-amz-fwd-header-Content-Range

The portion of the object returned in the response.

x-amz-fwd-header-Content-Type

A standard MIME type describing the format of the object data.

x-amz-fwd-header-ETag

An opaque identifier assigned by a web server to a specific version of a resource found at a URL.

x-amz-fwd-header-Expires

The date and time at which the object is no longer cacheable.

x-amz-fwd-header-Last-Modified

The date and time that the object was last modified.

x-amz-fwd-header-x-amz-checksum-crc32

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This specifies the Base64 encoded, 32-bit CRC32 checksum of the object returned by the Object Lambda function. This may not match the checksum for the object stored in Amazon S3. Amazon S3 will perform validation of the checksum values only

when the original `GetObject` request required checksum validation. For more information about checksums, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Only one checksum header can be specified at a time. If you supply multiple checksum headers, this request will fail.

x-amz-fwd-header-x-amz-checksum-crc32c

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This specifies the Base64 encoded, 32-bit CRC32C checksum of the object returned by the Object Lambda function. This may not match the checksum for the object stored in Amazon S3. Amazon S3 will perform validation of the checksum values only when the original `GetObject` request required checksum validation. For more information about checksums, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Only one checksum header can be specified at a time. If you supply multiple checksum headers, this request will fail.

x-amz-fwd-header-x-amz-checksum-crc64nvme

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 64-bit CRC64NVME checksum of the part. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-fwd-header-x-amz-checksum-md5

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 128-bit MD5 digest of the part. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-fwd-header-x-amz-checksum-sha1

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This specifies the Base64 encoded, 160-bit SHA1 digest of the object returned by the Object Lambda function. This may not match the checksum for the object stored in Amazon S3. Amazon S3 will perform validation of the checksum values only when the original `GetObject` request required checksum validation. For more information about checksums, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Only one checksum header can be specified at a time. If you supply multiple checksum headers, this request will fail.

x-amz-fwd-header-x-amz-checksum-sha256

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This specifies the Base64 encoded, 256-bit SHA256 digest of the object returned by the Object Lambda function. This may not match the checksum for the object stored in Amazon S3. Amazon S3 will perform validation of the checksum values only when the original `GetObject` request required checksum validation. For more information about checksums, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Only one checksum header can be specified at a time. If you supply multiple checksum headers, this request will fail.

x-amz-fwd-header-x-amz-checksum-sha512

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 512-bit SHA512 digest of the part. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-fwd-header-x-amz-checksum-xxhash128

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 128-bit XXHASH128 checksum of the part. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-fwd-header-x-amz-checksum-xxhash3

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 64-bit XXHASH3 checksum of the part. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-fwd-header-x-amz-checksum-xxhash64

This header can be used as a data integrity check to verify that the data received is the same data that was originally sent. This header specifies the Base64 encoded, 64-bit XXHASH64 checksum of the part. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

x-amz-fwd-header-x-amz-delete-marker

Specifies whether an object stored in Amazon S3 is (`true`) or is not (`false`) a delete marker. To learn more about delete markers, see [Working with delete markers](#).

[x-amz-fwd-header-x-amz-expiration](#)

If the object expiration is configured (see PUT Bucket lifecycle), the response includes this header. It includes the `expiry-date` and `rule-id` key-value pairs that provide the object expiration information. The value of the `rule-id` is URL-encoded.

[x-amz-fwd-header-x-amz-missing-meta](#)

Set to the number of metadata entries not returned in `x-amz-meta` headers. This can happen if you create metadata using an API like SOAP that supports more flexible metadata than the REST API. For example, using SOAP, you can create metadata whose values are not legal HTTP headers.

[x-amz-fwd-header-x-amz-mp-parts-count](#)

The count of parts this object has.

[x-amz-fwd-header-x-amz-object-lock-legal-hold](#)

Indicates whether an object stored in Amazon S3 has an active legal hold.

Valid Values: ON | OFF

[x-amz-fwd-header-x-amz-object-lock-mode](#)

Indicates whether an object stored in Amazon S3 has Object Lock enabled. For more information about S3 Object Lock, see [Object Lock](#).

Valid Values: GOVERNANCE | COMPLIANCE

[x-amz-fwd-header-x-amz-object-lock-retain-until-date](#)

The date and time when Object Lock is configured to expire.

[x-amz-fwd-header-x-amz-replication-status](#)

Indicates if request involves bucket that is either a source or destination in a Replication rule. For more information about S3 Replication, see [Replication](#).

Valid Values: COMPLETE | PENDING | FAILED | REPLICA | COMPLETED

[x-amz-fwd-header-x-amz-request-charged](#)

If present, indicates that the requester was successfully charged for the request. For more information, see [Using Requester Pays buckets for storage transfers and usage](#) in the *Amazon Simple Storage Service user guide*.

Note

This functionality is not supported for directory buckets.

Valid Values: `requester`

[x-amz-fwd-header-x-amz-restore](#)

Provides information about object restoration operation and expiration time of the restored object copy.

[x-amz-fwd-header-x-amz-server-side-encryption](#)

The server-side encryption algorithm used when storing requested object in Amazon S3 or Amazon FSx.

Note

When accessing data stored in Amazon FSx file systems using S3 access points, the only valid server side encryption option is `aws:fsx`.

Valid Values: `AES256` | `aws:fsx` | `aws:backup` | `aws:kms` | `aws:kms:dsse`

[x-amz-fwd-header-x-amz-server-side-encryption-aws-kms-key-id](#)

If present, specifies the ID (Key ID, Key ARN, or Key Alias) of the AWS Key Management Service (AWS KMS) symmetric encryption customer managed key that was used for stored in Amazon S3 object.

[x-amz-fwd-header-x-amz-server-side-encryption-bucket-key-enabled](#)

Indicates whether the object stored in Amazon S3 uses an S3 bucket key for server-side encryption with AWS KMS (SSE-KMS).

[x-amz-fwd-header-x-amz-server-side-encryption-customer-algorithm](#)

Encryption algorithm used if server-side encryption with a customer-provided encryption key was specified for object stored in Amazon S3.

x-amz-fwd-header-x-amz-server-side-encryption-customer-key-MD5

128-bit MD5 digest of customer-provided encryption key used in Amazon S3 to encrypt data stored in S3. For more information, see [Protecting data using server-side encryption with customer-provided encryption keys \(SSE-C\)](#).

x-amz-fwd-header-x-amz-storage-class

Provides storage class information of the object. Amazon S3 returns this header for all objects except for S3 Standard storage class objects.

For more information, see [Storage Classes](#).

Valid Values: STANDARD | REDUCED_REDUNDANCY | STANDARD_IA | ONEZONE_IA | INTELLIGENT_TIERING | GLACIER | DEEP_ARCHIVE | OUTPOSTS | GLACIER_IR | SNOW | EXPRESS_ONEZONE | FSX_OPENZFS | FSX_ONTAP | AWS_BACKUP_WARM | AWS_BACKUP_LOW_COST_WARM

x-amz-fwd-header-x-amz-tagging-count

The number of tags, if any, on the object.

x-amz-fwd-header-x-amz-version-id

An ID used to reference a specific version of the object.

x-amz-fwd-status

The integer status code for an HTTP response of a corresponding GetObject request. The following is a list of status codes.

- 200 - OK
- 206 - Partial Content
- 304 - Not Modified
- 400 - Bad Request
- 401 - Unauthorized
- 403 - Forbidden
- 404 - Not Found
- 405 - Method Not Allowed
- 409 - Conflict
- 411 - Length Required

- 412 - Precondition Failed
- 416 - Range Not Satisfiable
- 500 - Internal Server Error
- 503 - Service Unavailable

x-amz-request-route

Route prefix to the HTTP URL generated.

Required: Yes

x-amz-request-token

A single use encrypted token that maps WriteGetObjectResponse to the end user GetObject request.

Required: Yes

Request Body

The request accepts the following binary data.

Body

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Sample Response

The following illustrates a sample response.

```
HTTP/1.1 200 OK
```

```
x-amz-request-id: 19684529-d1aa-413e-9382-9ff490962d12
Date: Wed, 24 Feb 2021 10:57:53 GMT
Content-Length: 0
```

Sample Request

The following illustrates a sample request from a POST.

```
POST /WriteGetObjectResponse HTTP/1.1
Host: <RequestRoute>.s3-object-lambda.<Region>.amazonaws.com
x-amz-request-token: <RequestToken>
Authorization: authorization string
Content-Type: text/plain
Content-Length: 16
[16 bytes of object data]
```

Sample Error Response

The following response returns a `ValidationError` error because the `RequestToken` could not be decrypted.

```
<?xml version="1.0" encoding="UTF-8"?>
<Error>
  <Code>ValidationError</Code>
  <Message>Invalid token</Message>
  <RequestId>fcd2cd5e-def0-4001-8030-1fd1d61d2c9d</RequestId>
</Error>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)

- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

Amazon S3 Control

The following actions are supported by Amazon S3 Control:

- [AssociateAccessGrantsIdentityCenter](#)
- [CreateAccessGrant](#)
- [CreateAccessGrantsInstance](#)
- [CreateAccessGrantsLocation](#)
- [CreateAccessPoint](#)
- [CreateAccessPointForObjectLambda](#)
- [CreateBucket](#)
- [CreateJob](#)
- [CreateMultiRegionAccessPoint](#)
- [CreateStorageLensGroup](#)
- [DeleteAccessGrant](#)
- [DeleteAccessGrantsInstance](#)
- [DeleteAccessGrantsInstanceResourcePolicy](#)
- [DeleteAccessGrantsLocation](#)
- [DeleteAccessPoint](#)
- [DeleteAccessPointForObjectLambda](#)
- [DeleteAccessPointPolicy](#)
- [DeleteAccessPointPolicyForObjectLambda](#)
- [DeleteAccessPointScope](#)
- [DeleteBucket](#)
- [DeleteBucketLifecycleConfiguration](#)

- [DeleteBucketPolicy](#)
- [DeleteBucketReplication](#)
- [DeleteBucketTagging](#)
- [DeleteJobTagging](#)
- [DeleteMultiRegionAccessPoint](#)
- [DeletePublicAccessBlock](#)
- [DeleteStorageLensConfiguration](#)
- [DeleteStorageLensConfigurationTagging](#)
- [DeleteStorageLensGroup](#)
- [DescribeJob](#)
- [DescribeMultiRegionAccessPointOperation](#)
- [DissociateAccessGrantsIdentityCenter](#)
- [GetAccessGrant](#)
- [GetAccessGrantsInstance](#)
- [GetAccessGrantsInstanceForPrefix](#)
- [GetAccessGrantsInstanceResourcePolicy](#)
- [GetAccessGrantsLocation](#)
- [GetAccessPoint](#)
- [GetAccessPointConfigurationForObjectLambda](#)
- [GetAccessPointForObjectLambda](#)
- [GetAccessPointPolicy](#)
- [GetAccessPointPolicyForObjectLambda](#)
- [GetAccessPointPolicyStatus](#)
- [GetAccessPointPolicyStatusForObjectLambda](#)
- [GetAccessPointScope](#)
- [GetBucket](#)
- [GetBucketLifecycleConfiguration](#)
- [GetBucketPolicy](#)
- [GetBucketReplication](#)
- [GetBucketTagging](#)

- [GetBucketVersioning](#)
- [GetDataAccess](#)
- [GetJobTagging](#)
- [GetMultiRegionAccessPoint](#)
- [GetMultiRegionAccessPointPolicy](#)
- [GetMultiRegionAccessPointPolicyStatus](#)
- [GetMultiRegionAccessPointRoutes](#)
- [GetPublicAccessBlock](#)
- [GetStorageLensConfiguration](#)
- [GetStorageLensConfigurationTagging](#)
- [GetStorageLensGroup](#)
- [ListAccessGrants](#)
- [ListAccessGrantsInstances](#)
- [ListAccessGrantsLocations](#)
- [ListAccessPoints](#)
- [ListAccessPointsForDirectoryBuckets](#)
- [ListAccessPointsForObjectLambda](#)
- [ListCallerAccessGrants](#)
- [ListJobs](#)
- [ListMultiRegionAccessPoints](#)
- [ListRegionalBuckets](#)
- [ListStorageLensConfigurations](#)
- [ListStorageLensGroups](#)
- [ListTagsForResource](#)
- [PutAccessGrantsInstanceResourcePolicy](#)
- [PutAccessPointConfigurationForObjectLambda](#)
- [PutAccessPointPolicy](#)
- [PutAccessPointPolicyForObjectLambda](#)
- [PutAccessPointScope](#)
- [PutBucketLifecycleConfiguration](#)

- [PutBucketPolicy](#)
- [PutBucketReplication](#)
- [PutBucketTagging](#)
- [PutBucketVersioning](#)
- [PutJobTagging](#)
- [PutMultiRegionAccessPointPolicy](#)
- [PutPublicAccessBlock](#)
- [PutStorageLensConfiguration](#)
- [PutStorageLensConfigurationTagging](#)
- [SubmitMultiRegionAccessPointRoutes](#)
- [TagResource](#)
- [UntagResource](#)
- [UpdateAccessGrantsLocation](#)
- [UpdateJobPriority](#)
- [UpdateJobStatus](#)
- [UpdateStorageLensGroup](#)

AssociateAccessGrantsIdentityCenter

Service: Amazon S3 Control

Associate your S3 Access Grants instance with an AWS IAM Identity Center instance. Use this action if you want to create access grants for users or groups from your corporate identity directory. First, you must add your corporate identity directory to AWS IAM Identity Center. Then, you can associate this IAM Identity Center instance with your S3 Access Grants instance.

Permissions

You must have the `s3:AssociateAccessGrantsIdentityCenter` permission to use this operation.

Additional Permissions

You must also have the following permissions: `sso:CreateApplication`, `sso:PutApplicationGrant`, and `sso:PutApplicationAuthenticationMethod`.

Request Syntax

```
POST /v20180820/accessgrantsinstance/identitycenter HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
<?xml version="1.0" encoding="UTF-8"?>
<AssociateAccessGrantsIdentityCenterRequest xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <IdentityCenterArn>string</IdentityCenterArn>
</AssociateAccessGrantsIdentityCenterRequest>
```

URI Request Parameters

The request uses the following URI parameters.

[x-amz-account-id](#)

The AWS account ID of the S3 Access Grants instance.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request accepts the following data in XML format.

[AssociateAccessGrantsIdentityCenterRequest](#)

Root level tag for the AssociateAccessGrantsIdentityCenterRequest parameters.

Required: Yes

[IdentityCenterArn](#)

The Amazon Resource Name (ARN) of the AWS IAM Identity Center instance that you are associating with your S3 Access Grants instance. An IAM Identity Center instance is your corporate identity directory that you added to the IAM Identity Center. You can use the [ListInstances](#) API operation to retrieve a list of your Identity Center instances and their ARNs.

Type: String

Length Constraints: Minimum length of 10. Maximum length of 1224.

Pattern: `arn:[^:]+:sso:(\d{12}){0,1}:instance/.*$`

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)

- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

CreateAccessGrant

Service: Amazon S3 Control

Creates an access grant that gives a grantee access to your S3 data. The grantee can be an IAM user or role or a directory user, or group. Before you can create a grant, you must have an S3 Access Grants instance in the same Region as the S3 data. You can create an S3 Access Grants instance using the [CreateAccessGrantsInstance](#). You must also have registered at least one S3 data location in your S3 Access Grants instance using [CreateAccessGrantsLocation](#).

Permissions

You must have the `s3:CreateAccessGrant` permission to use this operation.

Additional Permissions

For any directory identity - `sso:DescribeInstance` and `sso:DescribeApplication`

For directory users - `identitystore:DescribeUser`

For directory groups - `identitystore:DescribeGroup`

Request Syntax

```
POST /v20180820/accessgrantsinstance/grant HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
<?xml version="1.0" encoding="UTF-8"?>
<CreateAccessGrantRequest xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <AccessGrantsLocationId>string</AccessGrantsLocationId>
  <AccessGrantsLocationConfiguration>
    <S3SubPrefix>string</S3SubPrefix>
  </AccessGrantsLocationConfiguration>
  <Grantee>
    <GranteeIdentifier>string</GranteeIdentifier>
    <GranteeType>string</GranteeType>
  </Grantee>
  <Permission>string</Permission>
  <ApplicationArn>string</ApplicationArn>
  <S3PrefixType>string</S3PrefixType>
  <Tags>
    <Tag>
      <Key>string</Key>
```

```
<Value>string</Value>
</Tag>
</Tags>
</CreateAccessGrantRequest>
```

URI Request Parameters

The request uses the following URI parameters.

x-amz-account-id

The AWS account ID of the S3 Access Grants instance.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request accepts the following data in XML format.

CreateAccessGrantRequest

Root level tag for the CreateAccessGrantRequest parameters.

Required: Yes

AccessGrantsLocationConfiguration

The configuration options of the grant location. The grant location is the S3 path to the data to which you are granting access. It contains the `S3SubPrefix` field. The grant scope is the result of appending the subprefix to the location scope of the registered location.

Type: [AccessGrantsLocationConfiguration](#) data type

Required: No

AccessGrantsLocationId

The ID of the registered location to which you are granting access. S3 Access Grants assigns this ID when you register the location. S3 Access Grants assigns the ID default to the default location `s3://` and assigns an auto-generated ID to other locations that you register.

If you are passing the default location, you cannot create an access grant for the entire default location. You must also specify a bucket or a bucket and prefix in the `Subprefix` field.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\_]+`

Required: Yes

ApplicationArn

The Amazon Resource Name (ARN) of an AWS IAM Identity Center application associated with your Identity Center instance. If an application ARN is included in the request to create an access grant, the grantee can only access the S3 data through this application.

Type: String

Length Constraints: Minimum length of 10. Maximum length of 1224.

Pattern: `arn:[^:]+\:sso::\d{12}:application/.*$`

Required: No

Grantee

The user, group, or role to which you are granting access. You can grant access to an IAM user or role. If you have added your corporate directory to AWS IAM Identity Center and associated your Identity Center instance with your S3 Access Grants instance, the grantee can also be a corporate directory user or group.

Type: [Grantee](#) data type

Required: Yes

Permission

The type of access that you are granting to your S3 data, which can be set to one of the following values:

- `READ` – Grant read-only access to the S3 data.
- `WRITE` – Grant write-only access to the S3 data.
- `READWRITE` – Grant both read and write access to the S3 data.

Type: String

Valid Values: READ | WRITE | READWRITE

Required: Yes

S3PrefixType

The type of S3SubPrefix. The only possible value is Object. Pass this value if the access grant scope is an object. Do not pass this value if the access grant scope is a bucket or a bucket and a prefix.

Type: String

Valid Values: Object

Required: No

Tags

The AWS resource tags that you are adding to the access grant. Each tag is a label consisting of a user-defined key and value. Tags can help you manage, identify, organize, search for, and filter resources.

Type: Array of [Tag](#) data types

Array Members: Minimum number of 0 items. Maximum number of 50 items.

Required: No

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<CreateAccessGrantResult>
  <CreatedAt>timestamp</CreatedAt>
  <AccessGrantId>string</AccessGrantId>
  <AccessGrantArn>string</AccessGrantArn>
  <Grantee>
    <GranteeIdentifier>string</GranteeIdentifier>
    <GranteeType>string</GranteeType>
  </Grantee>
  <AccessGrantsLocationId>string</AccessGrantsLocationId>
  <AccessGrantsLocationConfiguration>
```

```
<S3SubPrefix>string</S3SubPrefix>  
</AccessGrantsLocationConfiguration>  
<Permission>string</Permission>  
<ApplicationArn>string</ApplicationArn>  
<GrantScope>string</GrantScope>  
</CreateAccessGrantResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

CreateAccessGrantResult

Root level tag for the CreateAccessGrantResult parameters.

Required: Yes

AccessGrantArn

The Amazon Resource Name (ARN) of the access grant.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:[a-z\-\+]:s3:[a-z0-9\-\+]:\d{12}:access\-\grants\grant/[a-zA-Z0-9\-\+]`

AccessGrantId

The ID of the access grant. S3 Access Grants auto-generates this ID when you create the access grant.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\+]`

AccessGrantsLocationConfiguration

The configuration options of the grant location. The grant location is the S3 path to the data to which you are granting access.

Type: [AccessGrantsLocationConfiguration](#) data type

[AccessGrantsLocationId](#)

The ID of the registered location to which you are granting access. S3 Access Grants assigns this ID when you register the location. S3 Access Grants assigns the ID default to the default location `s3://` and assigns an auto-generated ID to other locations that you register.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\ ]+`

[ApplicationArn](#)

The Amazon Resource Name (ARN) of an AWS IAM Identity Center application associated with your Identity Center instance. If the grant includes an application ARN, the grantee can only access the S3 data through this application.

Type: String

Length Constraints: Minimum length of 10. Maximum length of 1224.

Pattern: `arn:[^:]+:sso::\d{12}:application/.*$`

[CreatedAt](#)

The date and time when you created the access grant.

Type: Timestamp

[Grantee](#)

The user, group, or role to which you are granting access. You can grant access to an IAM user or role. If you have added your corporate directory to AWS IAM Identity Center and associated your Identity Center instance with your S3 Access Grants instance, the grantee can also be a corporate directory user or group.

Type: [Grantee](#) data type

[GrantScope](#)

The S3 path of the data to which you are granting access. It is the result of appending the Subprefix to the location scope.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2000.

Pattern: `^\.+`

Permission

The type of access that you are granting to your S3 data, which can be set to one of the following values:

- READ – Grant read-only access to the S3 data.
- WRITE – Grant write-only access to the S3 data.
- READWRITE – Grant both read and write access to the S3 data.

Type: String

Valid Values: READ | WRITE | READWRITE

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

CreateAccessGrantsInstance

Service: Amazon S3 Control

Creates an S3 Access Grants instance, which serves as a logical grouping for access grants. You can create one S3 Access Grants instance per Region per account.

Permissions

You must have the `s3:CreateAccessGrantsInstance` permission to use this operation.

Additional Permissions

To associate an IAM Identity Center instance with your S3 Access Grants instance, you must also have the `sso:DescribeInstance`, `sso:CreateApplication`, `sso:PutApplicationGrant`, and `sso:PutApplicationAuthenticationMethod` permissions.

Request Syntax

```
POST /v20180820/accessgrantsinstance HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
<?xml version="1.0" encoding="UTF-8"?>
<CreateAccessGrantsInstanceRequest xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <IdentityCenterArn>string</IdentityCenterArn>
  <Tags>
    <Tag>
      <Key>string</Key>
      <Value>string</Value>
    </Tag>
  </Tags>
</CreateAccessGrantsInstanceRequest>
```

URI Request Parameters

The request uses the following URI parameters.

[x-amz-account-id](#)

The AWS account ID of the S3 Access Grants instance.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request accepts the following data in XML format.

CreateAccessGrantsInstanceRequest

Root level tag for the CreateAccessGrantsInstanceRequest parameters.

Required: Yes

IdentityCenterArn

If you would like to associate your S3 Access Grants instance with an AWS IAM Identity Center instance, use this field to pass the Amazon Resource Name (ARN) of the AWS IAM Identity Center instance that you are associating with your S3 Access Grants instance. An IAM Identity Center instance is your corporate identity directory that you added to the IAM Identity Center. You can use the [ListInstances](#) API operation to retrieve a list of your Identity Center instances and their ARNs.

Type: String

Length Constraints: Minimum length of 10. Maximum length of 1224.

Pattern: `arn:[^:]+:sso:(\d{12}){0,1}:instance/.*$`

Required: No

Tags

The AWS resource tags that you are adding to the S3 Access Grants instance. Each tag is a label consisting of a user-defined key and value. Tags can help you manage, identify, organize, search for, and filter resources.

Type: Array of [Tag](#) data types

Array Members: Minimum number of 0 items. Maximum number of 50 items.

Required: No

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<CreateAccessGrantsInstanceResult>
  <CreatedAt>timestamp</CreatedAt>
  <AccessGrantsInstanceId>string</AccessGrantsInstanceId>
  <AccessGrantsInstanceArn>string</AccessGrantsInstanceArn>
  <IdentityCenterArn>string</IdentityCenterArn>
  <IdentityCenterInstanceArn>string</IdentityCenterInstanceArn>
  <IdentityCenterApplicationArn>string</IdentityCenterApplicationArn>
</CreateAccessGrantsInstanceResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

CreateAccessGrantsInstanceResult

Root level tag for the CreateAccessGrantsInstanceResult parameters.

Required: Yes

AccessGrantsInstanceArn

The Amazon Resource Name (ARN) of the AWS IAM Identity Center instance that you are associating with your S3 Access Grants instance. An IAM Identity Center instance is your corporate identity directory that you added to the IAM Identity Center. You can use the [ListInstances](#) API operation to retrieve a list of your Identity Center instances and their ARNs.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:[a-z\-\+]:s3:[a-z0-9\-\+]:\d{12}:access\-\grants\[a-zA-Z0-9\-\+]`

AccessGrantsInstanceId

The ID of the S3 Access Grants instance. The ID is default. You can have one S3 Access Grants instance per Region per account.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [a-zA-Z0-9\-\-]+

CreatedAt

The date and time when you created the S3 Access Grants instance.

Type: Timestamp

IdentityCenterApplicationArn

If you associated your S3 Access Grants instance with an AWS IAM Identity Center instance, this field returns the Amazon Resource Name (ARN) of the IAM Identity Center instance application; a subresource of the original Identity Center instance. S3 Access Grants creates this Identity Center application for the specific S3 Access Grants instance.

Type: String

Length Constraints: Minimum length of 10. Maximum length of 1224.

Pattern: arn:[^:]+:sso:\d{12}:application/.*\$

IdentityCenterArn

This parameter has been deprecated.

If you associated your S3 Access Grants instance with an AWS IAM Identity Center instance, this field returns the Amazon Resource Name (ARN) of the IAM Identity Center instance application; a subresource of the original Identity Center instance. S3 Access Grants creates this Identity Center application for the specific S3 Access Grants instance.

Type: String

Length Constraints: Minimum length of 10. Maximum length of 1224.

Pattern: arn:[^:]+:sso:(\d{12}){0,1}:instance/.*\$

IdentityCenterInstanceArn

The Amazon Resource Name (ARN) of the AWS IAM Identity Center instance that you are associating with your S3 Access Grants instance. An IAM Identity Center instance is your corporate identity directory that you added to the IAM Identity Center. You can use the [ListInstances](#) API operation to retrieve a list of your Identity Center instances and their ARNs.

Type: String

Length Constraints: Minimum length of 10. Maximum length of 1224.

Pattern: `arn:[^:]+:sso:(\d{12}){0,1}:instance/.*$`

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

CreateAccessGrantsLocation

Service: Amazon S3 Control

The S3 data location that you would like to register in your S3 Access Grants instance. Your S3 data must be in the same Region as your S3 Access Grants instance. The location can be one of the following:

- The default S3 location `s3://`
- A bucket - `S3://<bucket-name>`
- A bucket and prefix - `S3://<bucket-name>/<prefix>`

When you register a location, you must include the IAM role that has permission to manage the S3 location that you are registering. Give S3 Access Grants permission to assume this role [using a policy](#). S3 Access Grants assumes this role to manage access to the location and to vend temporary credentials to grantees or client applications.

Permissions

You must have the `s3:CreateAccessGrantsLocation` permission to use this operation.

Additional Permissions

You must also have the following permission for the specified IAM role: `iam:PassRole`

Request Syntax

```
POST /v20180820/accessgrantsinstance/location HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
<?xml version="1.0" encoding="UTF-8"?>
<CreateAccessGrantsLocationRequest xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <LocationScope>string</LocationScope>
  <IAMRoleArn>string</IAMRoleArn>
  <Tags>
    <Tag>
      <Key>string</Key>
      <Value>string</Value>
    </Tag>
  </Tags>
```

```
</CreateAccessGrantsLocationRequest>
```

URI Request Parameters

The request uses the following URI parameters.

[x-amz-account-id](#)

The AWS account ID of the S3 Access Grants instance.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request accepts the following data in XML format.

[CreateAccessGrantsLocationRequest](#)

Root level tag for the CreateAccessGrantsLocationRequest parameters.

Required: Yes

[IAMRoleArn](#)

The Amazon Resource Name (ARN) of the IAM role for the registered location. S3 Access Grants assumes this role to manage access to the registered location.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:[^:]+:iam:\d{12}:role/.*`

Required: Yes

[LocationScope](#)

The S3 path to the location that you are registering. The location scope can be the default S3 location `s3://`, the S3 path to a bucket `s3://<bucket>`, or the S3 path to a bucket and prefix

s3://<bucket>/<prefix>. A prefix in S3 is a string of characters at the beginning of an object key name used to organize the objects that you store in your S3 buckets. For example, object key names that start with the `engineering/` prefix or object key names that start with the `marketing/campaigns/` prefix.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2000.

Pattern: `^\.+$`

Required: Yes

Tags

The AWS resource tags that you are adding to the S3 Access Grants location. Each tag is a label consisting of a user-defined key and value. Tags can help you manage, identify, organize, search for, and filter resources.

Type: Array of [Tag](#) data types

Array Members: Minimum number of 0 items. Maximum number of 50 items.

Required: No

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<CreateAccessGrantsLocationResult>
  <CreatedAt>timestamp</CreatedAt>
  <AccessGrantsLocationId>string</AccessGrantsLocationId>
  <AccessGrantsLocationArn>string</AccessGrantsLocationArn>
  <LocationScope>string</LocationScope>
  <IAMRoleArn>string</IAMRoleArn>
</CreateAccessGrantsLocationResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

CreateAccessGrantsLocationResult

Root level tag for the CreateAccessGrantsLocationResult parameters.

Required: Yes

AccessGrantsLocationArn

The Amazon Resource Name (ARN) of the location you are registering.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:[a-z\-\+]:s3:[a-z0-9\-\+]:\d{12}:access\-\grants\/location/[a-zA-Z0-9\-\+]`

AccessGrantsLocationId

The ID of the registered location to which you are granting access. S3 Access Grants assigns this ID when you register the location. S3 Access Grants assigns the ID default to the default location `s3://` and assigns an auto-generated ID to other locations that you register.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\+]`

CreatedAt

The date and time when you registered the location.

Type: Timestamp

IAMRoleArn

The Amazon Resource Name (ARN) of the IAM role for the registered location. S3 Access Grants assumes this role to manage access to the registered location.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:[^:]+\:iam::\d{12}:role/.*`

LocationScope

The S3 URI path to the location that you are registering. The location scope can be the default S3 location `s3://`, the S3 path to a bucket, or the S3 path to a bucket and prefix. A prefix in S3 is a string of characters at the beginning of an object key name used to organize the objects that you store in your S3 buckets. For example, object key names that start with the `engineering/` prefix or object key names that start with the `marketing/campaigns/` prefix.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2000.

Pattern: `^.+`

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

CreateAccessPoint

Service: Amazon S3 Control

Creates an access point and associates it to a specified bucket. For more information, see [Managing access to shared datasets with access points](#) or [Managing access to shared datasets in directory buckets with access points](#) in the *Amazon S3 User Guide*.

To create an access point and attach it to a volume on an Amazon FSx file system, see [CreateAndAttachS3AccessPoint](#) in the *Amazon FSx API Reference*.

Note

S3 on Outposts only supports VPC-style access points.

For more information, see [Accessing Amazon S3 on Outposts using virtual private cloud \(VPC\) only access points](#) in the *Amazon S3 User Guide*.

All Amazon S3 on Outposts REST API requests for this action require an additional parameter of `x-amz-outpost-id` to be passed with the request. In addition, you must use an S3 on Outposts endpoint hostname prefix instead of `s3-control`. For an example of the request syntax for Amazon S3 on Outposts that uses the S3 on Outposts endpoint hostname prefix and the `x-amz-outpost-id` derived by using the access point ARN, see the [Examples](#) section.

The following actions are related to `CreateAccessPoint`:

- [GetAccessPoint](#)
- [DeleteAccessPoint](#)
- [ListAccessPoints](#)
- [ListAccessPointsForDirectoryBuckets](#)

Request Syntax

```
PUT /v20180820/accesspoint/name HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
<?xml version="1.0" encoding="UTF-8"?>
<CreateAccessPointRequest xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
```

```

<Bucket>string</Bucket>
<VpcConfiguration>
  <VpcId>string</VpcId>
</VpcConfiguration>
<PublicAccessBlockConfiguration>
  <BlockPublicAcls>boolean</BlockPublicAcls>
  <BlockPublicPolicy>boolean</BlockPublicPolicy>
  <IgnorePublicAcls>boolean</IgnorePublicAcls>
  <RestrictPublicBuckets>boolean</RestrictPublicBuckets>
</PublicAccessBlockConfiguration>
<BucketAccountId>string</BucketAccountId>
<Scope>
  <Permissions>
    <Permission>string</Permission>
  </Permissions>
  <Prefixes>
    <Prefix>string</Prefix>
  </Prefixes>
</Scope>
<Tags>
  <Tag>
    <Key>string</Key>
    <Value>string</Value>
  </Tag>
</Tags>
</CreateAccessPointRequest>

```

URI Request Parameters

The request uses the following URI parameters.

name

The name you want to assign to this access point.

For directory buckets, the access point name must consist of a base name that you provide and suffix that includes the ZoneID (AWS Availability Zone or Local Zone) of your bucket location, followed by --xa-s3. For more information, see [Managing access to shared datasets in directory buckets with access points](#) in the *Amazon S3 User Guide*.

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

x-amz-account-id

The AWS account ID for the account that owns the specified access point.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request accepts the following data in XML format.

CreateAccessPointRequest

Root level tag for the CreateAccessPointRequest parameters.

Required: Yes

Bucket

The name of the bucket that you want to associate this access point with.

For using this parameter with Amazon S3 on Outposts with the REST API, you must specify the name and the x-amz-outpost-id as well.

For using this parameter with S3 on Outposts with the AWS SDK and CLI, you must specify the ARN of the bucket accessed in the format `arn:aws:s3-outposts:<Region>:<account-id>:outpost/<outpost-id>/bucket/<my-bucket-name>`. For example, to access the bucket `reports` through Outpost `my-outpost` owned by account `123456789012` in Region `us-west-2`, use the URL encoding of `arn:aws:s3-outposts:us-west-2:123456789012:outpost/my-outpost/bucket/reports`. The value must be URL encoded.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

BucketAccountId

The AWS account ID associated with the S3 bucket associated with this access point.

For same account access point when your bucket and access point belong to the same account owner, the `BucketAccountId` is not required. For cross-account access point when your bucket and access point are not in the same account, the `BucketAccountId` is required.

Type: String

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: No

PublicAccessBlockConfiguration

The `PublicAccessBlock` configuration that you want to apply to the access point.

Type: [PublicAccessBlockConfiguration](#) data type

Required: No

Scope

For directory buckets, you can filter access control to specific prefixes, API operations, or a combination of both. For more information, see [Managing access to shared datasets in directory buckets with access points](#) in the *Amazon S3 User Guide*.

Note

Scope is only supported for access points attached to directory buckets.

Type: [Scope](#) data type

Required: No

Tags

An array of tags that you can apply to an access point. Tags are key-value pairs of metadata used to control access to your access points. For more information about tags, see [Using tags with Amazon S3](#). For information about tagging access points, see [Using tags for attribute-based access control \(ABAC\)](#).

Type: Array of [Tag](#) data types

Array Members: Minimum number of 0 items. Maximum number of 50 items.

Required: No

VpcConfiguration

If you include this field, Amazon S3 restricts access to this access point to requests from the specified virtual private cloud (VPC).

Note

This is required for creating an access point for Amazon S3 on Outposts buckets.

Type: [VpcConfiguration](#) data type

Required: No

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<CreateAccessPointResult>
  <AccessPointArn>string</AccessPointArn>
  <Alias>string</Alias>
</CreateAccessPointResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

CreateAccessPointResult

Root level tag for the CreateAccessPointResult parameters.

Required: Yes

AccessPointArn

The ARN of the access point.

Note

This is only supported by Amazon S3 on Outposts.

Type: String

Length Constraints: Minimum length of 4. Maximum length of 128.

Alias

The name or alias of the access point.

Type: String

Length Constraints: Maximum length of 63.

Pattern: `^[0-9a-z\\-]{63}`

Examples**Sample request for creating an access point for an Amazon S3 on Outposts bucket**

This request creates an access point for S3 on Outposts bucket.

```
PUT /v20180820/accesspoint/example-access-point HTTP/1.1
Host:s3-outposts.<Region>.amazonaws.com
x-amz-account-id: example-account-id
x-amz-outpost-id: op-01ac5d28a6a232904
<?xml version="1.0" encoding="UTF-8"?>
  <CreateAccessPointRequest xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
    <Bucket>example-outpost-bucket </Bucket>
  </CreateAccessPointRequest>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

CreateAccessPointForObjectLambda

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Creates an Object Lambda Access Point. For more information, see [Transforming objects with Object Lambda Access Points](#) in the *Amazon S3 User Guide*.

The following actions are related to CreateAccessPointForObjectLambda:

- [DeleteAccessPointForObjectLambda](#)
- [GetAccessPointForObjectLambda](#)
- [ListAccessPointsForObjectLambda](#)

Request Syntax

```
PUT /v20180820/accesspointforobjectlambda/name HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
<?xml version="1.0" encoding="UTF-8"?>
<CreateAccessPointForObjectLambdaRequest xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <Configuration>
    <AllowedFeatures>
      <AllowedFeature>string</AllowedFeature>
    </AllowedFeatures>
    <CloudWatchMetricsEnabled>boolean</CloudWatchMetricsEnabled>
    <SupportingAccessPoint>string</SupportingAccessPoint>
    <TransformationConfigurations>
      <TransformationConfiguration>
        <Actions>
          <Action>string</Action>
        </Actions>
        <ContentTransformation>
          <AwsLambda>
            <FunctionArn>string</FunctionArn>
            <FunctionPayload>string</FunctionPayload>
```

```
</AwsLambda>
</ContentTransformation>
</TransformationConfiguration>
</TransformationConfigurations>
</Configuration>
</CreateAccessPointForObjectLambdaRequest>
```

URI Request Parameters

The request uses the following URI parameters.

name

The name you want to assign to this Object Lambda Access Point.

Length Constraints: Minimum length of 3. Maximum length of 45.

Pattern: `^[a-z0-9]([a-z0-9\-*][a-z0-9])?$`

Required: Yes

x-amz-account-id

The AWS account ID for owner of the specified Object Lambda Access Point.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request accepts the following data in XML format.

CreateAccessPointForObjectLambdaRequest

Root level tag for the CreateAccessPointForObjectLambdaRequest parameters.

Required: Yes

Configuration

Object Lambda Access Point configuration as a JSON document.

Type: [ObjectLambdaConfiguration](#) data type

Required: Yes

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<CreateAccessPointForObjectLambdaResult>
  <ObjectLambdaAccessPointArn>string</ObjectLambdaAccessPointArn>
  <Alias>
    <Status>string</Status>
    <Value>string</Value>
  </Alias>
</CreateAccessPointForObjectLambdaResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[CreateAccessPointForObjectLambdaResult](#)

Root level tag for the CreateAccessPointForObjectLambdaResult parameters.

Required: Yes

[Alias](#)

The alias of the Object Lambda Access Point.

Type: [ObjectLambdaAccessPointAlias](#) data type

[ObjectLambdaAccessPointArn](#)

Specifies the ARN for the Object Lambda Access Point.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:[^:]+:s3-object-lambda:[^:]*:\d{12}:accesspoint/.*`

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

CreateBucket

Service: Amazon S3 Control

Note

This action creates an Amazon S3 on Outposts bucket. To create an S3 bucket, see [Create Bucket](#) in the *Amazon S3 API Reference*.

Creates a new Outposts bucket. By creating the bucket, you become the bucket owner. To create an Outposts bucket, you must have S3 on Outposts. For more information, see [Using Amazon S3 on Outposts](#) in *Amazon S3 User Guide*.

Not every string is an acceptable bucket name. For information on bucket naming restrictions, see [Working with Amazon S3 Buckets](#).

S3 on Outposts buckets support:

- Tags
- LifecycleConfigurations for deleting expired objects

For a complete list of restrictions and Amazon S3 feature limitations on S3 on Outposts, see [Amazon S3 on Outposts Restrictions and Limitations](#).

For an example of the request syntax for Amazon S3 on Outposts that uses the S3 on Outposts endpoint hostname prefix and `x-amz-outpost-id` in your API request, see the [Examples](#) section.

The following actions are related to CreateBucket for Amazon S3 on Outposts:

- [PutObject](#)
- [GetBucket](#)
- [DeleteBucket](#)
- [CreateAccessPoint](#)
- [PutAccessPointPolicy](#)

Request Syntax

```
PUT /v20180820/bucket/name HTTP/1.1
```

```
Host: Bucket.s3-control.amazonaws.com
x-amz-acl: ACL
x-amz-grant-full-control: GrantFullControl
x-amz-grant-read: GrantRead
x-amz-grant-read-acp: GrantReadACP
x-amz-grant-write: GrantWrite
x-amz-grant-write-acp: GrantWriteACP
x-amz-bucket-object-lock-enabled: ObjectLockEnabledForBucket
x-amz-outpost-id: OutpostId
<?xml version="1.0" encoding="UTF-8"?>
<CreateBucketConfiguration xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <LocationConstraint>string</LocationConstraint>
</CreateBucketConfiguration>
```

URI Request Parameters

The request uses the following URI parameters.

[name](#)

The name of the bucket.

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

[x-amz-acl](#)

The canned ACL to apply to the bucket.

Note

This is not supported by Amazon S3 on Outposts buckets.

Valid Values: private | public-read | public-read-write | authenticated-read

[x-amz-bucket-object-lock-enabled](#)

Specifies whether you want S3 Object Lock to be enabled for the new bucket.

Note

This is not supported by Amazon S3 on Outposts buckets.

x-amz-grant-full-control

Allows grantee the read, write, read ACP, and write ACP permissions on the bucket.

 Note

This is not supported by Amazon S3 on Outposts buckets.

x-amz-grant-read

Allows grantee to list the objects in the bucket.

 Note

This is not supported by Amazon S3 on Outposts buckets.

x-amz-grant-read-acp

Allows grantee to read the bucket ACL.

 Note

This is not supported by Amazon S3 on Outposts buckets.

x-amz-grant-write

Allows grantee to create, overwrite, and delete any object in the bucket.

 Note

This is not supported by Amazon S3 on Outposts buckets.

x-amz-grant-write-acp

Allows grantee to write the ACL for the applicable bucket.

Note

This is not supported by Amazon S3 on Outposts buckets.

x-amz-outpost-id

The ID of the Outposts where the bucket is being created.

Note

This ID is required by Amazon S3 on Outposts buckets.

Length Constraints: Minimum length of 1. Maximum length of 64.

Request Body

The request accepts the following data in XML format.

CreateBucketConfiguration

Root level tag for the CreateBucketConfiguration parameters.

Required: Yes

LocationConstraint

Specifies the Region where the bucket will be created. If you are creating a bucket on the US East (N. Virginia) Region (us-east-1), you do not need to specify the location.

Note

This is not supported by Amazon S3 on Outposts buckets.

Type: String

Valid Values: EU | eu-west-1 | us-west-1 | us-west-2 | ap-south-1 | ap-southeast-1 | ap-southeast-2 | ap-northeast-1 | sa-east-1 | cn-north-1 | eu-central-1

Required: No

Response Syntax

```
HTTP/1.1 200
Location: Location
<?xml version="1.0" encoding="UTF-8"?>
<CreateBucketResult>
  <BucketArn>string</BucketArn>
</CreateBucketResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The response returns the following HTTP headers.

Location

The location of the bucket.

The following data is returned in XML format by the service.

CreateBucketResult

Root level tag for the CreateBucketResult parameters.

Required: Yes

BucketArn

The Amazon Resource Name (ARN) of the bucket.

For using this parameter with Amazon S3 on Outposts with the REST API, you must specify the name and the x-amz-outpost-id as well.

For using this parameter with S3 on Outposts with the AWS SDK and CLI, you must specify the ARN of the bucket accessed in the format `arn:aws:s3-outposts:<Region>:<account-id>:outpost/<outpost-id>/bucket/<my-bucket-name>`. For example, to access the bucket `reports` through Outpost `my-outpost` owned by account `123456789012` in Region `us-west-2`, use the URL encoding of `arn:aws:s3-outposts:us-`

west-2:123456789012:outpost/my-outpost/bucket/reports. The value must be URL encoded.

Type: String

Length Constraints: Minimum length of 4. Maximum length of 128.

Errors

BucketAlreadyExists

The requested Outposts bucket name is not available. The bucket namespace is shared by all users of the AWS Outposts in this Region. Select a different name and try again.

HTTP Status Code: 400

BucketAlreadyOwnedByYou

The Outposts bucket you tried to create already exists, and you own it.

HTTP Status Code: 400

Examples

Sample request to create an Amazon S3 on Outposts bucket

This request creates an Outposts bucket named `example-outpost-bucket`.

```
PUT /v20180820/bucket/example-outpost-bucket/ HTTP/1.1
Host:s3-outposts.<Region>.amazonaws.com
x-amz-outpost-id: op-01ac5d28a6a232904
Content-Length:
Date: Wed, 01 Mar 2006 12:00:00 GMT
Authorization: authorization string
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

CreateJob

Service: Amazon S3 Control

This operation creates an S3 Batch Operations job.

You can use S3 Batch Operations to perform large-scale batch actions on Amazon S3 objects. Batch Operations can run a single action on lists of Amazon S3 objects that you specify. For more information, see [S3 Batch Operations](#) in the *Amazon S3 User Guide*.

Permissions

For information about permissions required to use the Batch Operations, see [Granting permissions for S3 Batch Operations](#) in the *Amazon S3 User Guide*.

Related actions include:

- [DescribeJob](#)
- [ListJobs](#)
- [UpdateJobPriority](#)
- [UpdateJobStatus](#)
- [JobOperation](#)

Request Syntax

```
POST /v20180820/jobs HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
<?xml version="1.0" encoding="UTF-8"?>
<CreateJobRequest xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <ConfirmationRequired>boolean</ConfirmationRequired>
  <Operation>
    <LambdaInvoke>
      <FunctionArn>string</FunctionArn>
      <InvocationSchemaVersion>string</InvocationSchemaVersion>
      <UserArguments>
        <entry>
          <key>string</key>
          <value>string</value>
        </entry>
```

```

    </UserArguments>
  </LambdaInvoke>
  <S3ComputeObjectChecksum>
    <ChecksumAlgorithm>string</ChecksumAlgorithm>
    <ChecksumType>string</ChecksumType>
  </S3ComputeObjectChecksum>
  <S3DeleteObjectTagging>
  </S3DeleteObjectTagging>
  <S3InitiateRestoreObject>
    <ExpirationInDays>integer</ExpirationInDays>
    <GlacierJobTier>string</GlacierJobTier>
  </S3InitiateRestoreObject>
  <S3PutObjectAcl>
    <AccessControlPolicy>
      <AccessControlList>
        <Grants>
          <S3Grant>
            <Grantee>
              <DisplayName>string</DisplayName>
              <Identifier>string</Identifier>
              <TypeIdentifier>string</TypeIdentifier>
            </Grantee>
            <Permission>string</Permission>
          </S3Grant>
        </Grants>
        <Owner>
          <DisplayName>string</DisplayName>
          <ID>string</ID>
        </Owner>
      </AccessControlList>
      <CannedAccessControlList>string</CannedAccessControlList>
    </AccessControlPolicy>
  </S3PutObjectAcl>
  <S3PutObjectCopy>
    <AccessControlGrants>
      <S3Grant>
        <Grantee>
          <DisplayName>string</DisplayName>
          <Identifier>string</Identifier>
          <TypeIdentifier>string</TypeIdentifier>
        </Grantee>
        <Permission>string</Permission>
      </S3Grant>
    </AccessControlGrants>
  </S3PutObjectCopy>

```

```

<AnnotationDirective>string</AnnotationDirective>
<BucketKeyEnabled>boolean</BucketKeyEnabled>
<CannedAccessControlList>string</CannedAccessControlList>
<ChecksumAlgorithm>string</ChecksumAlgorithm>
<MetadataDirective>string</MetadataDirective>
<ModifiedSinceConstraint>timestamp</ModifiedSinceConstraint>
<NewObjectMetadata>
  <CacheControl>string</CacheControl>
  <ContentDisposition>string</ContentDisposition>
  <ContentEncoding>string</ContentEncoding>
  <ContentLanguage>string</ContentLanguage>
  <ContentLength>long</ContentLength>
  <ContentMD5>string</ContentMD5>
  <ContentType>string</ContentType>
  <HttpExpiresDate>timestamp</HttpExpiresDate>
  <RequesterCharged>boolean</RequesterCharged>
  <SSEAlgorithm>string</SSEAlgorithm>
  <UserMetadata>
    <entry>
      <key>string</key>
      <value>string</value>
    </entry>
  </UserMetadata>
</NewObjectMetadata>
<NewObjectTagging>
  <S3Tag>
    <Key>string</Key>
    <Value>string</Value>
  </S3Tag>
</NewObjectTagging>
<ObjectLockEventHold>string</ObjectLockEventHold>
<ObjectLockEventHoldDuration>
  <Days>integer</Days>
  <Years>integer</Years>
</ObjectLockEventHoldDuration>
<ObjectLockLegalHoldStatus>string</ObjectLockLegalHoldStatus>
<ObjectLockMode>string</ObjectLockMode>
<ObjectLockRetainUntilDate>timestamp</ObjectLockRetainUntilDate>
<RedirectLocation>string</RedirectLocation>
<RequesterPays>boolean</RequesterPays>
<SSEAwsKmsKeyId>string</SSEAwsKmsKeyId>
<StorageClass>string</StorageClass>
<TargetKeyPrefix>string</TargetKeyPrefix>
<TargetResource>string</TargetResource>

```



```

    <UnmodifiedSinceConstraint>timestamp</UnmodifiedSinceConstraint>
  </S3PutObjectCopy>
  <S3PutObjectLegalHold>
    <LegalHold>
      <Status>string</Status>
    </LegalHold>
  </S3PutObjectLegalHold>
  <S3PutObjectRetention>
    <BypassGovernanceRetention>boolean</BypassGovernanceRetention>
    <Retention>
      <EventHold>string</EventHold>
      <EventHoldDuration>
        <Days>integer</Days>
        <Years>integer</Years>
      </EventHoldDuration>
      <Mode>string</Mode>
      <RetainUntilDate>timestamp</RetainUntilDate>
    </Retention>
  </S3PutObjectRetention>
  <S3PutObjectTagging>
    <TagSet>
      <S3Tag>
        <Key>string</Key>
        <Value>string</Value>
      </S3Tag>
    </TagSet>
  </S3PutObjectTagging>
  <S3ReplicateObject>
  </S3ReplicateObject>
  <S3UpdateObjectEncryption>
    <ObjectEncryption>
      <SSE-KMS>
        <BucketKeyEnabled>boolean</BucketKeyEnabled>
        <KMSKeyArn>string</KMSKeyArn>
      </SSE-KMS>
    </ObjectEncryption>
  </S3UpdateObjectEncryption>
</Operation>
<Report>
  <Bucket>string</Bucket>
  <Enabled>boolean</Enabled>
  <ExpectedBucketOwner>string</ExpectedBucketOwner>
  <Format>string</Format>
  <Prefix>string</Prefix>

```

```

    <ReportScope>string</ReportScope>
  </Report>
  <ClientRequestToken>string</ClientRequestToken>
  <Manifest>
    <Location>
      <ETag>string</ETag>
      <ObjectArn>string</ObjectArn>
      <ObjectVersionId>string</ObjectVersionId>
    </Location>
    <Spec>
      <Fields>
        <member>string</member>
      </Fields>
      <Format>string</Format>
    </Spec>
  </Manifest>
  <Description>string</Description>
  <Priority>integer</Priority>
  <RoleArn>string</RoleArn>
  <Tags>
    <S3Tag>
      <Key>string</Key>
      <Value>string</Value>
    </S3Tag>
  </Tags>
  <ManifestGenerator>
    <S3JobManifestGenerator>
      <EnableManifestOutput>boolean</EnableManifestOutput>
      <ExpectedBucketOwner>string</ExpectedBucketOwner>
      <Filter>
        <CreatedAfter>timestamp</CreatedAfter>
        <CreatedBefore>timestamp</CreatedBefore>
        <EligibleForReplication>boolean</EligibleForReplication>
        <KeyNameConstraint>
          <MatchAnyPrefix>
            <member>string</member>
          </MatchAnyPrefix>
          <MatchAnySubstring>
            <member>string</member>
          </MatchAnySubstring>
          <MatchAnySuffix>
            <member>string</member>
          </MatchAnySuffix>
        </KeyNameConstraint>
      </Filter>
    </S3JobManifestGenerator>
  </ManifestGenerator>

```

```

<MatchAnyObjectEncryption>
  <ObjectEncryption>
    <DSSE-KMS>
      <KmsKeyArn>string</KmsKeyArn>
    </DSSE-KMS>
    <NOT-SSE>
    </NOT-SSE>
    <SSE-C>
    </SSE-C>
    <SSE-KMS>
      <BucketKeyEnabled>boolean</BucketKeyEnabled>
      <KmsKeyArn>string</KmsKeyArn>
    </SSE-KMS>
    <SSE-S3>
    </SSE-S3>
  </ObjectEncryption>
</MatchAnyObjectEncryption>
<MatchAnyStorageClass>
  <member>string</member>
</MatchAnyStorageClass>
<ObjectReplicationStatuses>
  <member>string</member>
</ObjectReplicationStatuses>
<ObjectSizeGreaterThanBytes>long</ObjectSizeGreaterThanBytes>
<ObjectSizeLessThanBytes>long</ObjectSizeLessThanBytes>
</Filter>
<ManifestOutputLocation>
  <Bucket>string</Bucket>
  <ExpectedManifestBucketOwner>string</ExpectedManifestBucketOwner>
  <ManifestEncryption>
    <SSE-KMS>
      <KeyId>string</KeyId>
    </SSE-KMS>
    <SSE-S3>
    </SSE-S3>
  </ManifestEncryption>
  <ManifestFormat>string</ManifestFormat>
  <ManifestPrefix>string</ManifestPrefix>
</ManifestOutputLocation>
<SourceBucket>string</SourceBucket>
</S3JobManifestGenerator>
</ManifestGenerator>
</CreateJobRequest>

```

URI Request Parameters

The request uses the following URI parameters.

[x-amz-account-id](#)

The AWS account ID that creates the job.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request accepts the following data in XML format.

[CreateJobRequest](#)

Root level tag for the CreateJobRequest parameters.

Required: Yes

[ClientRequestToken](#)

An idempotency token to ensure that you don't accidentally submit the same request twice. You can use any string up to the maximum length.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Required: Yes

[ConfirmationRequired](#)

Indicates whether confirmation is required before Amazon S3 runs the job. Confirmation is only required for jobs created through the Amazon S3 console.

Type: Boolean

Required: No

Description

A description for this job. You can use any string within the permitted length. Descriptions don't need to be unique and can be used for multiple jobs.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Required: No

Manifest

Configuration parameters for the manifest.

Type: [JobManifest](#) data type

Required: No

ManifestGenerator

The attribute container for the ManifestGenerator details. Jobs must be created with either a manifest file or a ManifestGenerator, but not both.

Type: [JobManifestGenerator](#) data type

Note: This object is a Union. Only one member of this object can be specified or returned.

Required: No

Operation

The action that you want this job to perform on every object listed in the manifest. For more information about the available actions, see [Operations](#) in the *Amazon S3 User Guide*.

Type: [JobOperation](#) data type

Required: Yes

Priority

The numerical priority for this job. Higher numbers indicate higher priority.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 2147483647.

Required: Yes

Report

Configuration parameters for the optional job-completion report.

Type: [JobReport](#) data type

Required: Yes

RoleArn

The Amazon Resource Name (ARN) for the AWS Identity and Access Management (IAM) role that Batch Operations will use to run this job's action on every object in the manifest.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:[^:]+:iam::\d{12}:role/.*`

Required: Yes

Tags

A set of tags to associate with the S3 Batch Operations job. This is an optional parameter.

Type: Array of [S3Tag](#) data types

Required: No

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<CreateJobResult>
  <JobId>string</JobId>
</CreateJobResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

CreateJobResult

Root level tag for the CreateJobResult parameters.

Required: Yes

JobId

The ID for this job. Amazon S3 generates this ID automatically and returns it after a successful Create Job request.

Type: String

Length Constraints: Minimum length of 5. Maximum length of 36.

Pattern: [a-zA-Z0-9\-__]+

Errors

BadRequestException

HTTP Status Code: 400

IdempotencyException

HTTP Status Code: 400

InternalServiceException

HTTP Status Code: 500

TooManyRequestsException

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)

- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

CreateMultiRegionAccessPoint

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Creates a Multi-Region Access Point and associates it with the specified buckets. For more information about creating Multi-Region Access Points, see [Creating Multi-Region Access Points](#) in the *Amazon S3 User Guide*.

This action will always be routed to the US West (Oregon) Region. For more information about the restrictions around working with Multi-Region Access Points, see [Multi-Region Access Point restrictions and limitations](#) in the *Amazon S3 User Guide*.

This request is asynchronous, meaning that you might receive a response before the command has completed. When this request provides a response, it provides a token that you can use to monitor the status of the request with `DescribeMultiRegionAccessPointOperation`.

The following actions are related to `CreateMultiRegionAccessPoint`:

- [DeleteMultiRegionAccessPoint](#)
- [DescribeMultiRegionAccessPointOperation](#)
- [GetMultiRegionAccessPoint](#)
- [ListMultiRegionAccessPoints](#)

Request Syntax

```
POST /v20180820/async-requests/mrap/create HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
<?xml version="1.0" encoding="UTF-8"?>
<CreateMultiRegionAccessPointRequest xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <ClientToken>string</ClientToken>
  <Details>
    <Name>string</Name>
    <PublicAccessBlock>
```

```
<BlockPublicAcls>boolean</BlockPublicAcls>
<BlockPublicPolicy>boolean</BlockPublicPolicy>
<IgnorePublicAcls>boolean</IgnorePublicAcls>
<RestrictPublicBuckets>boolean</RestrictPublicBuckets>
</PublicAccessBlock>
<Regions>
  <Region>
    <Bucket>string</Bucket>
    <BucketAccountId>string</BucketAccountId>
  </Region>
</Regions>
</Details>
</CreateMultiRegionAccessPointRequest>
```

URI Request Parameters

The request uses the following URI parameters.

[x-amz-account-id](#)

The AWS account ID for the owner of the Multi-Region Access Point. The owner of the Multi-Region Access Point also must own the underlying buckets.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request accepts the following data in XML format.

[CreateMultiRegionAccessPointRequest](#)

Root level tag for the CreateMultiRegionAccessPointRequest parameters.

Required: Yes

[ClientToken](#)

An idempotency token used to identify the request and guarantee that requests are unique.

Type: String

Length Constraints: Maximum length of 64.

Pattern: \S+

Required: Yes

Details

A container element containing details about the Multi-Region Access Point.

Type: [CreateMultiRegionAccessPointInput](#) data type

Required: Yes

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<CreateMultiRegionAccessPointResult>
  <RequestTokenARN>string</RequestTokenARN>
</CreateMultiRegionAccessPointResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

CreateMultiRegionAccessPointResult

Root level tag for the CreateMultiRegionAccessPointResult parameters.

Required: Yes

RequestTokenARN

The request token associated with the request. You can use this token with [DescribeMultiRegionAccessPointOperation](#) to determine the status of asynchronous requests.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `arn: . +`

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

CreateStorageLensGroup

Service: Amazon S3 Control

Creates a new S3 Storage Lens group and associates it with the specified AWS account ID. An S3 Storage Lens group is a custom grouping of objects based on prefix, suffix, object tags, object size, object age, or a combination of these filters. For each Storage Lens group that you've created, you can also optionally add AWS resource tags. For more information about S3 Storage Lens groups, see [Working with S3 Storage Lens groups](#).

To use this operation, you must have the permission to perform the `s3:CreateStorageLensGroup` action. If you're trying to create a Storage Lens group with AWS resource tags, you must also have permission to perform the `s3:TagResource` action. For more information about the required Storage Lens Groups permissions, see [Setting account permissions to use S3 Storage Lens groups](#).

For information about Storage Lens groups errors, see [List of Amazon S3 Storage Lens error codes](#).

Request Syntax

```
POST /v20180820/storagegroup HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
<?xml version="1.0" encoding="UTF-8"?>
<CreateStorageLensGroupRequest xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <StorageLensGroup>
    <Filter>
      <And>
        <MatchAnyPrefix>
          <Prefix>string</Prefix>
        </MatchAnyPrefix>
        <MatchAnySuffix>
          <Suffix>string</Suffix>
        </MatchAnySuffix>
        <MatchAnyTag>
          <Tag>
            <Keystring</Key>
            <Valuestring</Value>
          </Tag>
        </MatchAnyTag>
        <MatchObjectAge>
          <DaysGreaterThanOrEqualTointeger</DaysGreaterThanOrEqualTo>
```

```

    <DaysLessThan>integer</DaysLessThan>
  </MatchObjectAge>
  <MatchObjectSize>
    <BytesGreaterThan>long</BytesGreaterThan>
    <BytesLessThan>long</BytesLessThan>
  </MatchObjectSize>
</And>
<MatchAnyPrefix>
  <Prefix>string</Prefix>
</MatchAnyPrefix>
<MatchAnySuffix>
  <Suffix>string</Suffix>
</MatchAnySuffix>
<MatchAnyTag>
  <Tag>
    <Key>string</Key>
    <Value>string</Value>
  </Tag>
</MatchAnyTag>
<MatchObjectAge>
  <DaysGreaterThan>integer</DaysGreaterThan>
  <DaysLessThan>integer</DaysLessThan>
</MatchObjectAge>
<MatchObjectSize>
  <BytesGreaterThan>long</BytesGreaterThan>
  <BytesLessThan>long</BytesLessThan>
</MatchObjectSize>
<Or>
  <MatchAnyPrefix>
    <Prefix>string</Prefix>
  </MatchAnyPrefix>
  <MatchAnySuffix>
    <Suffix>string</Suffix>
  </MatchAnySuffix>
  <MatchAnyTag>
    <Tag>
      <Key>string</Key>
      <Value>string</Value>
    </Tag>
  </MatchAnyTag>
  <MatchObjectAge>
    <DaysGreaterThan>integer</DaysGreaterThan>
    <DaysLessThan>integer</DaysLessThan>
  </MatchObjectAge>

```

```

    <MatchObjectSize>
      <BytesGreaterThan>long</BytesGreaterThan>
      <BytesLessThan>long</BytesLessThan>
    </MatchObjectSize>
  </Or>
</Filter>
<Name>string</Name>
<StorageLensGroupArn>string</StorageLensGroupArn>
</StorageLensGroup>
<Tags>
  <Tag>
    <Key>string</Key>
    <Value>string</Value>
  </Tag>
</Tags>
</CreateStorageLensGroupRequest>

```

URI Request Parameters

The request uses the following URI parameters.

x-amz-account-id

The AWS account ID that the Storage Lens group is created from and associated with.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request accepts the following data in XML format.

CreateStorageLensGroupRequest

Root level tag for the CreateStorageLensGroupRequest parameters.

Required: Yes

StorageLensGroup

The Storage Lens group configuration.

Type: [StorageLensGroup](#) data type

Required: Yes

Tags

The AWS resource tags that you're adding to your Storage Lens group. This parameter is optional.

Type: Array of [Tag](#) data types

Array Members: Minimum number of 0 items. Maximum number of 50 items.

Required: No

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteAccessGrant

Service: Amazon S3 Control

Deletes the access grant from the S3 Access Grants instance. You cannot undo an access grant deletion and the grantee will no longer have access to the S3 data.

Permissions

You must have the `s3:DeleteAccessGrant` permission to use this operation.

Request Syntax

```
DELETE /v20180820/accessgrantsinstance/grant/id HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

id

The ID of the access grant. S3 Access Grants auto-generates this ID when you create the access grant.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\_]+`

Required: Yes

x-amz-account-id

The AWS account ID of the S3 Access Grants instance.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteAccessGrantsInstance

Service: Amazon S3 Control

Deletes your S3 Access Grants instance. You must first delete the access grants and locations before S3 Access Grants can delete the instance. See [DeleteAccessGrant](#) and [DeleteAccessGrantsLocation](#). If you have associated an IAM Identity Center instance with your S3 Access Grants instance, you must first disassociate the Identity Center instance from the S3 Access Grants instance before you can delete the S3 Access Grants instance. See [AssociateAccessGrantsIdentityCenter](#) and [DissociateAccessGrantsIdentityCenter](#).

Permissions

You must have the `s3:DeleteAccessGrantsInstance` permission to use this operation.

Request Syntax

```
DELETE /v20180820/accessgrantsinstance HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

[x-amz-account-id](#)

The AWS account ID of the S3 Access Grants instance.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteAccessGrantsInstanceResourcePolicy

Service: Amazon S3 Control

Deletes the resource policy of the S3 Access Grants instance. The resource policy is used to manage cross-account access to your S3 Access Grants instance. By deleting the resource policy, you delete any cross-account permissions to your S3 Access Grants instance.

Permissions

You must have the `s3:DeleteAccessGrantsInstanceResourcePolicy` permission to use this operation.

Request Syntax

```
DELETE /v20180820/accessgrantsinstance/resourcepolicy HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

x-amz-account-id

The AWS account ID of the S3 Access Grants instance.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteAccessGrantsLocation

Service: Amazon S3 Control

Deregisters a location from your S3 Access Grants instance. You can only delete a location registration from an S3 Access Grants instance if there are no grants associated with this location. See [Delete a grant](#) for information on how to delete grants. You need to have at least one registered location in your S3 Access Grants instance in order to create access grants.

Permissions

You must have the `s3:DeleteAccessGrantsLocation` permission to use this operation.

Request Syntax

```
DELETE /v20180820/accessgrantsinstance/location/id HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

id

The ID of the registered location that you are deregistering from your S3 Access Grants instance. S3 Access Grants assigned this ID when you registered the location. S3 Access Grants assigns the ID default to the default location `s3://` and assigns an auto-generated ID to other locations that you register.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\_]+`

Required: Yes

x-amz-account-id

The AWS account ID of the S3 Access Grants instance.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteAccessPoint

Service: Amazon S3 Control

Deletes the specified access point.

All Amazon S3 on Outposts REST API requests for this action require an additional parameter of `x-amz-outpost-id` to be passed with the request. In addition, you must use an S3 on Outposts endpoint hostname prefix instead of `s3-control`. For an example of the request syntax for Amazon S3 on Outposts that uses the S3 on Outposts endpoint hostname prefix and the `x-amz-outpost-id` derived by using the access point ARN, see the [Examples](#) section.

The following actions are related to `DeleteAccessPoint`:

- [CreateAccessPoint](#)
- [GetAccessPoint](#)
- [ListAccessPoints](#)

Request Syntax

```
DELETE /v20180820/accesspoint/name HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

[name](#)

The name of the access point you want to delete.

For using this parameter with Amazon S3 on Outposts with the REST API, you must specify the name and the `x-amz-outpost-id` as well.

For using this parameter with S3 on Outposts with the AWS SDK and CLI, you must specify the ARN of the access point accessed in the format `arn:aws:s3-outposts:<Region>:<account-id>:outpost/<outpost-id>/accesspoint/<my-accesspoint-name>`. For example, to access the access point `reports-ap` through Outpost `my-outpost` owned by account `123456789012` in Region `us-west-2`, use the URL

encoding of `arn:aws:s3-outposts:us-west-2:123456789012:outpost/my-outpost/accesspoint/reports-ap`. The value must be URL encoded.

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

x-amz-account-id

The AWS account ID for the account that owns the specified access point.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

DeleteAccessPoint syntax for Amazon S3 on Outposts

The following request deletes the access point of the specified Outpost.

```
DELETE /v20180820/accesspoint/example-access-point HTTP/1.1
Host: s3-outposts.<Region>.amazonaws.com
Date: Wed, 28 Oct 2020 22:32:00 GMT
x-amz-account-id: example-account-id
x-amz-outpost-id: op-01ac5d28a6a232904
Authorization: authorization string
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteAccessPointForObjectLambda

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Deletes the specified Object Lambda Access Point.

The following actions are related to DeleteAccessPointForObjectLambda:

- [CreateAccessPointForObjectLambda](#)
- [GetAccessPointForObjectLambda](#)
- [ListAccessPointsForObjectLambda](#)

Request Syntax

```
DELETE /v20180820/accesspointforobjectlambda/name HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the access point you want to delete.

Length Constraints: Minimum length of 3. Maximum length of 45.

Pattern: `^[a-z0-9]([a-z0-9\-*][a-z0-9])?$`

Required: Yes

x-amz-account-id

The account ID for the account that owns the specified Object Lambda Access Point.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteAccessPointPolicy

Service: Amazon S3 Control

Deletes the access point policy for the specified access point.

All Amazon S3 on Outposts REST API requests for this action require an additional parameter of `x-amz-outpost-id` to be passed with the request. In addition, you must use an S3 on Outposts endpoint hostname prefix instead of `s3-control`. For an example of the request syntax for Amazon S3 on Outposts that uses the S3 on Outposts endpoint hostname prefix and the `x-amz-outpost-id` derived by using the access point ARN, see the [Examples](#) section.

The following actions are related to `DeleteAccessPointPolicy`:

- [PutAccessPointPolicy](#)
- [GetAccessPointPolicy](#)

Request Syntax

```
DELETE /v20180820/accesspoint/name/policy HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

[name](#)

The name of the access point whose policy you want to delete.

For using this parameter with Amazon S3 on Outposts with the REST API, you must specify the name and the `x-amz-outpost-id` as well.

For using this parameter with S3 on Outposts with the AWS SDK and CLI, you must specify the ARN of the access point accessed in the format `arn:aws:s3-outposts:<Region>:<account-id>:outpost/<outpost-id>/accesspoint/<my-accesspoint-name>`. For example, to access the access point `reports-ap` through Outpost `my-outpost` owned by account `123456789012` in Region `us-west-2`, use the URL

encoding of `arn:aws:s3-outposts:us-west-2:123456789012:outpost/my-outpost/accesspoint/reports-ap`. The value must be URL encoded.

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

x-amz-account-id

The account ID for the account that owns the specified access point.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Sample request syntax for using the DeleteAccessPointPolicy action with Amazon S3 on Outposts access point

This example illustrates one usage of DeleteAccessPointPolicy.

```
DELETE /v20180820/accesspoint/example-access-point/policy HTTP/1.1
Host: s3-outposts.<Region>.amazonaws.com
Date: Wed, 28 Oct 2020 22:32:00 GMT
Authorization: authorization string
x-amz-account-id: example-account-id
```



```
x-amz-outpost-id: op-01ac5d28a6a232904
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteAccessPointPolicyForObjectLambda

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Removes the resource policy for an Object Lambda Access Point.

The following actions are related to DeleteAccessPointPolicyForObjectLambda:

- [GetAccessPointPolicyForObjectLambda](#)
- [PutAccessPointPolicyForObjectLambda](#)

Request Syntax

```
DELETE /v20180820/accesspointforobjectlambda/name/policy HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the Object Lambda Access Point you want to delete the policy for.

Length Constraints: Minimum length of 3. Maximum length of 45.

Pattern: `^[a-z0-9]([a-z0-9\-*][a-z0-9])?$`

Required: Yes

x-amz-account-id

The account ID for the account that owns the specified Object Lambda Access Point.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteAccessPointScope

Service: Amazon S3 Control

Deletes an existing access point scope for a directory bucket.

Note

When you delete the scope of an access point, all prefixes and permissions are deleted.

To use this operation, you must have the permission to perform the `s3express:DeleteAccessPointScope` action.

For information about REST API errors, see [REST error responses](#).

Request Syntax

```
DELETE /v20180820/accesspoint/name/scope HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the access point with the scope that you want to delete.

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

x-amz-account-id

The AWS account ID that owns the access point with the scope that you want to delete.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteBucket

Service: Amazon S3 Control

Note

This action deletes an Amazon S3 on Outposts bucket. To delete an S3 bucket, see [DeleteBucket](#) in the *Amazon S3 API Reference*.

Deletes the Amazon S3 on Outposts bucket. All objects (including all object versions and delete markers) in the bucket must be deleted before the bucket itself can be deleted. For more information, see [Using Amazon S3 on Outposts](#) in *Amazon S3 User Guide*.

All Amazon S3 on Outposts REST API requests for this action require an additional parameter of `x-amz-outpost-id` to be passed with the request. In addition, you must use an S3 on Outposts endpoint hostname prefix instead of `s3-control`. For an example of the request syntax for Amazon S3 on Outposts that uses the S3 on Outposts endpoint hostname prefix and the `x-amz-outpost-id` derived by using the access point ARN, see the [Examples](#) section.

Related Resources

- [CreateBucket](#)
- [GetBucket](#)
- [DeleteObject](#)

Request Syntax

```
DELETE /v20180820/bucket/name HTTP/1.1
Host: Bucket.s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

name

Specifies the bucket being deleted.

For using this parameter with Amazon S3 on Outposts with the REST API, you must specify the name and the `x-amz-outpost-id` as well.

For using this parameter with S3 on Outposts with the AWS SDK and CLI, you must specify the ARN of the bucket accessed in the format `arn:aws:s3-outposts:<Region>:<account-id>:outpost/<outpost-id>/bucket/<my-bucket-name>`. For example, to access the bucket `reports` through Outpost `my-outpost` owned by account `123456789012` in Region `us-west-2`, use the URL encoding of `arn:aws:s3-outposts:us-west-2:123456789012:outpost/my-outpost/bucket/reports`. The value must be URL encoded.

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

`x-amz-account-id`

The account ID that owns the Outposts bucket.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Sample request to delete an Amazon S3 on Outposts bucket

This request deletes the Outposts bucket named `example-outpost-bucket`.

```
DELETE /v20180820/bucket/example-outpost-bucket/ HTTP/1.1
Host: s3-outposts.<Region>.amazonaws.com
x-amz-outpost-id: op-01ac5d28a6a232904
x-amz-account-id:example-account-id
Date: Wed, 01 Mar 2006 12:00:00 GMT
Authorization: authorization string
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteBucketLifecycleConfiguration

Service: Amazon S3 Control

Note

This action deletes an Amazon S3 on Outposts bucket's lifecycle configuration. To delete an S3 bucket's lifecycle configuration, see [DeleteBucketLifecycle](#) in the *Amazon S3 API Reference*.

Deletes the lifecycle configuration from the specified Outposts bucket. Amazon S3 on Outposts removes all the lifecycle configuration rules in the lifecycle subresource associated with the bucket. Your objects never expire, and Amazon S3 on Outposts no longer automatically deletes any objects on the basis of rules contained in the deleted lifecycle configuration. For more information, see [Using Amazon S3 on Outposts](#) in *Amazon S3 User Guide*.

To use this operation, you must have permission to perform the `s3-outposts:PutLifecycleConfiguration` action. By default, the bucket owner has this permission and the Outposts bucket owner can grant this permission to others.

All Amazon S3 on Outposts REST API requests for this action require an additional parameter of `x-amz-outpost-id` to be passed with the request. In addition, you must use an S3 on Outposts endpoint hostname prefix instead of `s3-control`. For an example of the request syntax for Amazon S3 on Outposts that uses the S3 on Outposts endpoint hostname prefix and the `x-amz-outpost-id` derived by using the access point ARN, see the [Examples](#) section.

For more information about object expiration, see [Elements to Describe Lifecycle Actions](#).

Related actions include:

- [PutBucketLifecycleConfiguration](#)
- [GetBucketLifecycleConfiguration](#)

Request Syntax

```
DELETE /v20180820/bucket/name/lifecycleconfiguration HTTP/1.1
Host: Bucket.s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

name

Specifies the bucket.

For using this parameter with Amazon S3 on Outposts with the REST API, you must specify the name and the x-amz-outpost-id as well.

For using this parameter with S3 on Outposts with the AWS SDK and CLI, you must specify the ARN of the bucket accessed in the format `arn:aws:s3-outposts:<Region>:<account-id>:outpost/<outpost-id>/bucket/<my-bucket-name>`. For example, to access the bucket `reports` through Outpost `my-outpost` owned by account `123456789012` in Region `us-west-2`, use the URL encoding of `arn:aws:s3-outposts:us-west-2:123456789012:outpost/my-outpost/bucket/reports`. The value must be URL encoded.

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

x-amz-account-id

The account ID of the lifecycle configuration to delete.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Sample request to delete the lifecycle configuration of an Amazon S3 on Outposts bucket

This example illustrates one usage of DeleteBucketLifecycleConfiguration.

```
DELETE /v20180820/bucket/example-outpost-bucket/
lifecycleconfiguration HTTP/1.1
Host: s3-outposts.<Region>.amazonaws.com
x-amz-outpost-id: op-01ac5d28a6a232904
x-amz-account-id:example-account-id
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteBucketPolicy

Service: Amazon S3 Control

Note

This action deletes an Amazon S3 on Outposts bucket policy. To delete an S3 bucket policy, see [DeleteBucketPolicy](#) in the *Amazon S3 API Reference*.

This implementation of the DELETE action uses the policy subresource to delete the policy of a specified Amazon S3 on Outposts bucket. If you are using an identity other than the root user of the AWS account that owns the bucket, the calling identity must have the `s3-outposts:DeleteBucketPolicy` permissions on the specified Outposts bucket and belong to the bucket owner's account to use this action. For more information, see [Using Amazon S3 on Outposts](#) in *Amazon S3 User Guide*.

If you don't have `DeleteBucketPolicy` permissions, Amazon S3 returns a `403 Access Denied` error. If you have the correct permissions, but you're not using an identity that belongs to the bucket owner's account, Amazon S3 returns a `405 Method Not Allowed` error.

Important

As a security precaution, the root user of the AWS account that owns a bucket can always use this action, even if the policy explicitly denies the root user the ability to perform this action.

For more information about bucket policies, see [Using Bucket Policies and User Policies](#).

All Amazon S3 on Outposts REST API requests for this action require an additional parameter of `x-amz-outpost-id` to be passed with the request. In addition, you must use an S3 on Outposts endpoint hostname prefix instead of `s3-control`. For an example of the request syntax for Amazon S3 on Outposts that uses the S3 on Outposts endpoint hostname prefix and the `x-amz-outpost-id` derived by using the access point ARN, see the [Examples](#) section.

The following actions are related to `DeleteBucketPolicy`:

- [GetBucketPolicy](#)

- [PutBucketPolicy](#)

Request Syntax

```
DELETE /v20180820/bucket/name/policy HTTP/1.1
Host: Bucket.s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

[name](#)

Specifies the bucket.

For using this parameter with Amazon S3 on Outposts with the REST API, you must specify the name and the x-amz-outpost-id as well.

For using this parameter with S3 on Outposts with the AWS SDK and CLI, you must specify the ARN of the bucket accessed in the format `arn:aws:s3-outposts:<Region>:<account-id>:outpost/<outpost-id>/bucket/<my-bucket-name>`. For example, to access the bucket `reports` through Outpost `my-outpost` owned by account `123456789012` in Region `us-west-2`, use the URL encoding of `arn:aws:s3-outposts:us-west-2:123456789012:outpost/my-outpost/bucket/reports`. The value must be URL encoded.

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

[x-amz-account-id](#)

The account ID of the Outposts bucket.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Sample request for deleting a bucket policy for an Amazon S3 on Outposts bucket

This example illustrates one usage of DeleteBucketPolicy.

```
DELETE v20180820/bucket/example-outpost-bucket/policy HTTP/1.1
Host: s3-outposts.<Region>.amazonaws.com
x-amz-account-id: example-account-id
x-amz-outpost-id: op-01ac5d28a6a232904
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)

- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteBucketReplication

Service: Amazon S3 Control

Note

This operation deletes an Amazon S3 on Outposts bucket's replication configuration. To delete an S3 bucket's replication configuration, see [DeleteBucketReplication](#) in the *Amazon S3 API Reference*.

Deletes the replication configuration from the specified S3 on Outposts bucket.

To use this operation, you must have permissions to perform the `s3-outposts:PutReplicationConfiguration` action. The Outposts bucket owner has this permission by default and can grant it to others. For more information about permissions, see [Setting up IAM with S3 on Outposts](#) and [Managing access to S3 on Outposts buckets](#) in the *Amazon S3 User Guide*.

Note

It can take a while to propagate PUT or DELETE requests for a replication configuration to all S3 on Outposts systems. Therefore, the replication configuration that's returned by a GET request soon after a PUT or DELETE request might return a more recent result than what's on the Outpost. If an Outpost is offline, the delay in updating the replication configuration on that Outpost can be significant.

All Amazon S3 on Outposts REST API requests for this action require an additional parameter of `x-amz-outpost-id` to be passed with the request. In addition, you must use an S3 on Outposts endpoint hostname prefix instead of `s3-control`. For an example of the request syntax for Amazon S3 on Outposts that uses the S3 on Outposts endpoint hostname prefix and the `x-amz-outpost-id` derived by using the access point ARN, see the [Examples](#) section.

For information about S3 replication on Outposts configuration, see [Replicating objects for S3 on Outposts](#) in the *Amazon S3 User Guide*.

The following operations are related to `DeleteBucketReplication`:

- [PutBucketReplication](#)

- [GetBucketReplication](#)

Request Syntax

```
DELETE /v20180820/bucket/name/replication HTTP/1.1
Host: Bucket.s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

[name](#)

Specifies the S3 on Outposts bucket to delete the replication configuration for.

For using this parameter with Amazon S3 on Outposts with the REST API, you must specify the name and the x-amz-outpost-id as well.

For using this parameter with S3 on Outposts with the AWS SDK and CLI, you must specify the ARN of the bucket accessed in the format `arn:aws:s3-outposts:<Region>:<account-id>:outpost/<outpost-id>/bucket/<my-bucket-name>`. For example, to access the bucket `reports` through Outpost `my-outpost` owned by account `123456789012` in Region `us-west-2`, use the URL encoding of `arn:aws:s3-outposts:us-west-2:123456789012:outpost/my-outpost/bucket/reports`. The value must be URL encoded.

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

[x-amz-account-id](#)

The AWS account ID of the Outposts bucket to delete the replication configuration for.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Sample Request

The following DELETE request deletes the replication subresource from the specified S3 on Outposts bucket. This request removes the replication configuration that is set for the bucket.

```
DELETE /v20180820/bucket/example-outpost-bucket/replication HTTP/1.1
Host: s3-outposts.<Region>.amazonaws.com
x-amz-outpost-id: op-01ac5d28a6a232904
x-amz-account-id:example-account-id
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)

- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteBucketTagging

Service: Amazon S3 Control

Note

This action deletes an Amazon S3 on Outposts bucket's tags. To delete an S3 bucket tags, see [DeleteBucketTagging](#) in the *Amazon S3 API Reference*.

Deletes the tags from the Outposts bucket. For more information, see [Using Amazon S3 on Outposts](#) in *Amazon S3 User Guide*.

To use this action, you must have permission to perform the PutBucketTagging action. By default, the bucket owner has this permission and can grant this permission to others.

All Amazon S3 on Outposts REST API requests for this action require an additional parameter of `x-amz-outpost-id` to be passed with the request. In addition, you must use an S3 on Outposts endpoint hostname prefix instead of `s3-control`. For an example of the request syntax for Amazon S3 on Outposts that uses the S3 on Outposts endpoint hostname prefix and the `x-amz-outpost-id` derived by using the access point ARN, see the [Examples](#) section.

The following actions are related to DeleteBucketTagging:

- [GetBucketTagging](#)
- [PutBucketTagging](#)

Request Syntax

```
DELETE /v20180820/bucket/name/tagging HTTP/1.1
Host: Bucket.s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

name

The bucket ARN that has the tag set to be removed.

For using this parameter with Amazon S3 on Outposts with the REST API, you must specify the name and the `x-amz-outpost-id` as well.

For using this parameter with S3 on Outposts with the AWS SDK and CLI, you must specify the ARN of the bucket accessed in the format `arn:aws:s3-outposts:<Region>:<account-id>:outpost/<outpost-id>/bucket/<my-bucket-name>`. For example, to access the bucket `reports` through Outpost `my-outpost` owned by account `123456789012` in Region `us-west-2`, use the URL encoding of `arn:aws:s3-outposts:us-west-2:123456789012:outpost/my-outpost/bucket/reports`. The value must be URL encoded.

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

`x-amz-account-id`

The AWS account ID of the Outposts bucket tag set to be removed.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Examples

Sample request to delete tags for Amazon S3 on Outposts bucket

The following DELETE request deletes the tag set from the Outposts bucket `example-outpost-bucket`.

```
DELETE v20180820/bucket/example-outpost-bucket/tagging HTTP/1.1
Host: s3-outposts.<Region>.amazonaws.com
x-amz-account-id: example-account-id
x-amz-outpost-id: op-01ac5d28a6a232904
Date: Wed, 14 Dec 2020 05:37:16 GMT
Authorization: signatureValue
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteJobTagging

Service: Amazon S3 Control

Removes the entire tag set from the specified S3 Batch Operations job.

Permissions

To use the DeleteJobTagging operation, you must have permission to perform the `s3:DeleteJobTagging` action. For more information, see [Controlling access and labeling jobs using tags](#) in the *Amazon S3 User Guide*.

Related actions include:

- [CreateJob](#)
- [GetJobTagging](#)
- [PutJobTagging](#)

Request Syntax

```
DELETE /v20180820/jobs/id/tagging HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

id

The ID for the S3 Batch Operations job whose tags you want to delete.

Length Constraints: Minimum length of 5. Maximum length of 36.

Pattern: `[a-zA-Z0-9\-\_\~]+`

Required: Yes

x-amz-account-id

The AWS account ID associated with the S3 Batch Operations job.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

InternalServerError

HTTP Status Code: 500

NotFoundException

HTTP Status Code: 400

TooManyRequestsException

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)

- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteMultiRegionAccessPoint

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Deletes a Multi-Region Access Point. This action does not delete the buckets associated with the Multi-Region Access Point, only the Multi-Region Access Point itself.

This action will always be routed to the US West (Oregon) Region. For more information about the restrictions around working with Multi-Region Access Points, see [Multi-Region Access Point restrictions and limitations](#) in the *Amazon S3 User Guide*.

This request is asynchronous, meaning that you might receive a response before the command has completed. When this request provides a response, it provides a token that you can use to monitor the status of the request with `DescribeMultiRegionAccessPointOperation`.

The following actions are related to `DeleteMultiRegionAccessPoint`:

- [CreateMultiRegionAccessPoint](#)
- [DescribeMultiRegionAccessPointOperation](#)
- [GetMultiRegionAccessPoint](#)
- [ListMultiRegionAccessPoints](#)

Request Syntax

```
POST /v20180820/async-requests/mrap/delete HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
<?xml version="1.0" encoding="UTF-8"?>
<DeleteMultiRegionAccessPointRequest xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <ClientToken>string</ClientToken>
  <Details>
    <Name>string</Name>
  </Details>
</DeleteMultiRegionAccessPointRequest>
```

URI Request Parameters

The request uses the following URI parameters.

[x-amz-account-id](#)

The AWS account ID for the owner of the Multi-Region Access Point.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request accepts the following data in XML format.

[DeleteMultiRegionAccessPointRequest](#)

Root level tag for the DeleteMultiRegionAccessPointRequest parameters.

Required: Yes

[ClientToken](#)

An idempotency token used to identify the request and guarantee that requests are unique.

Type: String

Length Constraints: Maximum length of 64.

Pattern: `\S+`

Required: Yes

[Details](#)

A container element containing details about the Multi-Region Access Point.

Type: [DeleteMultiRegionAccessPointInput](#) data type

Required: Yes

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<DeleteMultiRegionAccessPointResult>
  <RequestTokenARN>string</RequestTokenARN>
</DeleteMultiRegionAccessPointResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[DeleteMultiRegionAccessPointResult](#)

Root level tag for the DeleteMultiRegionAccessPointResult parameters.

Required: Yes

[RequestTokenARN](#)

The request token associated with the request. You can use this token with [DescribeMultiRegionAccessPointOperation](#) to determine the status of asynchronous requests.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `arn: . +`

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeletePublicAccessBlock

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Removes the `PublicAccessBlock` configuration for an AWS account. This operation might be restricted when the account is managed by organization-level Block Public Access policies. You'll get an Access Denied (403) error when the account is managed by organization-level Block Public Access policies. Organization-level policies override account-level settings, preventing direct account-level modifications. For more information, see [Using Amazon S3 block public access](#).

Related actions include:

- [GetPublicAccessBlock](#)
- [PutPublicAccessBlock](#)

Request Syntax

```
DELETE /v20180820/configuration/publicAccessBlock HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

[x-amz-account-id](#)

The account ID for the AWS account whose `PublicAccessBlock` configuration you want to remove.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteStorageLensConfiguration

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Deletes the Amazon S3 Storage Lens configuration. For more information about S3 Storage Lens, see [Assessing your storage activity and usage with Amazon S3 Storage Lens](#) in the *Amazon S3 User Guide*.

Note

To use this action, you must have permission to perform the `s3:DeleteStorageLensConfiguration` action. For more information, see [Setting permissions to use Amazon S3 Storage Lens](#) in the *Amazon S3 User Guide*.

Request Syntax

```
DELETE /v20180820/storagelens/storagelensid HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

storagelensid

The ID of the S3 Storage Lens configuration.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\_\.\ ]+`

Required: Yes

x-amz-account-id

The account ID of the requester.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteStorageLensConfigurationTagging

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Deletes the Amazon S3 Storage Lens configuration tags. For more information about S3 Storage Lens, see [Assessing your storage activity and usage with Amazon S3 Storage Lens](#) in the *Amazon S3 User Guide*.

Note

To use this action, you must have permission to perform the `s3:DeleteStorageLensConfigurationTagging` action. For more information, see [Setting permissions to use Amazon S3 Storage Lens](#) in the *Amazon S3 User Guide*.

Request Syntax

```
DELETE /v20180820/storagelens/storage lens id/tagging HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

storage lens id

The ID of the S3 Storage Lens configuration.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\_\.\ ]+`

Required: Yes

x-amz-account-id

The account ID of the requester.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteStorageLensGroup

Service: Amazon S3 Control

Deletes an existing S3 Storage Lens group.

To use this operation, you must have the permission to perform the `s3:DeleteStorageLensGroup` action. For more information about the required Storage Lens Groups permissions, see [Setting account permissions to use S3 Storage Lens groups](#).

For information about Storage Lens groups errors, see [List of Amazon S3 Storage Lens error codes](#).

Request Syntax

```
DELETE /v20180820/storagelensgroup/name HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the Storage Lens group that you're trying to delete.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\_\_]+`

Required: Yes

x-amz-account-id

The AWS account ID used to create the Storage Lens group that you're trying to delete.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DescribeJob

Service: Amazon S3 Control

Retrieves the configuration parameters and status for a Batch Operations job. For more information, see [S3 Batch Operations](#) in the *Amazon S3 User Guide*.

Permissions

To use the DescribeJob operation, you must have permission to perform the `s3:DescribeJob` action.

Related actions include:

- [CreateJob](#)
- [ListJobs](#)
- [UpdateJobPriority](#)
- [UpdateJobStatus](#)

Request Syntax

```
GET /v20180820/jobs/id HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

id

The ID for the job whose information you want to retrieve.

Length Constraints: Minimum length of 5. Maximum length of 36.

Pattern: `[a-zA-Z0-9\-\_\_]+`

Required: Yes

x-amz-account-id

The AWS account ID associated with the S3 Batch Operations job.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<DescribeJobResult>
  <Job>
    <ConfirmationRequired>boolean</ConfirmationRequired>
    <CreationTime>timestamp</CreationTime>
    <Description>string</Description>
    <FailureReasons>
      <JobFailure>
        <FailureCode>string</FailureCode>
        <FailureReason>string</FailureReason>
      </JobFailure>
    </FailureReasons>
    <GeneratedManifestDescriptor>
      <Format>string</Format>
      <Location>
        <ETag>string</ETag>
        <ObjectArn>string</ObjectArn>
        <ObjectVersionId>string</ObjectVersionId>
      </Location>
    </GeneratedManifestDescriptor>
    <JobArn>string</JobArn>
    <JobId>string</JobId>
    <Manifest>
      <Location>
        <ETag>string</ETag>
        <ObjectArn>string</ObjectArn>
        <ObjectVersionId>string</ObjectVersionId>
      </Location>
    <Spec>
      <Fields>
```

```

        <member>string</member>
    </Fields>
    <Format>string</Format>
</Spec>
</Manifest>
<ManifestGenerator>
    <S3JobManifestGenerator>
        <EnableManifestOutput>boolean</EnableManifestOutput>
        <ExpectedBucketOwner>string</ExpectedBucketOwner>
        <Filter>
            <CreatedAfter>timestamp</CreatedAfter>
            <CreatedBefore>timestamp</CreatedBefore>
            <EligibleForReplication>boolean</EligibleForReplication>
            <KeyNameConstraint>
                <MatchAnyPrefix>
                    <member>string</member>
                </MatchAnyPrefix>
                <MatchAnySubstring>
                    <member>string</member>
                </MatchAnySubstring>
                <MatchAnySuffix>
                    <member>string</member>
                </MatchAnySuffix>
            </KeyNameConstraint>
            <MatchAnyObjectEncryption>
                <ObjectEncryption>
                    <DSSE-KMS>
                        <KmsKeyArn>string</KmsKeyArn>
                    </DSSE-KMS>
                    <NOT-SSE>
                    </NOT-SSE>
                    <SSE-C>
                    </SSE-C>
                    <SSE-KMS>
                        <BucketKeyEnabled>boolean</BucketKeyEnabled>
                        <KmsKeyArn>string</KmsKeyArn>
                    </SSE-KMS>
                    <SSE-S3>
                    </SSE-S3>
                </ObjectEncryption>
            </MatchAnyObjectEncryption>
            <MatchAnyStorageClass>
                <member>string</member>
            </MatchAnyStorageClass>
        </ManifestGenerator>
    </S3JobManifestGenerator>
</ManifestGenerator>

```

```

    <ObjectReplicationStatuses>
      <member>string</member>
    </ObjectReplicationStatuses>
    <ObjectSizeGreaterThanBytes>long</ObjectSizeGreaterThanBytes>
    <ObjectSizeLessThanBytes>long</ObjectSizeLessThanBytes>
  </Filter>
  <ManifestOutputLocation>
    <Bucket>string</Bucket>
    <ExpectedManifestBucketOwner>string</ExpectedManifestBucketOwner>
    <ManifestEncryption>
      <SSE-KMS>
        <KeyId>string</KeyId>
      </SSE-KMS>
      <SSE-S3>
      </SSE-S3>
    </ManifestEncryption>
    <ManifestFormat>string</ManifestFormat>
    <ManifestPrefix>string</ManifestPrefix>
  </ManifestOutputLocation>
  <SourceBucket>string</SourceBucket>
</S3JobManifestGenerator>
</ManifestGenerator>
<Operation>
  <LambdaInvoke>
    <FunctionArn>string</FunctionArn>
    <InvocationSchemaVersion>string</InvocationSchemaVersion>
    <UserArguments>
      <entry>
        <key>string</key>
        <value>string</value>
      </entry>
    </UserArguments>
  </LambdaInvoke>
  <S3ComputeObjectChecksum>
    <ChecksumAlgorithm>string</ChecksumAlgorithm>
    <ChecksumType>string</ChecksumType>
  </S3ComputeObjectChecksum>
  <S3DeleteObjectTagging>
  </S3DeleteObjectTagging>
  <S3InitiateRestoreObject>
    <ExpirationInDays>integer</ExpirationInDays>
    <GlacierJobTier>string</GlacierJobTier>
  </S3InitiateRestoreObject>
  <S3PutObjectAcl>

```

```

    <AccessControlPolicy>
      <AccessControlList>
        <Grants>
          <S3Grant>
            <Grantee>
              <DisplayName>string</DisplayName>
              <Identifier>string</Identifier>
              <TypeIdentifier>string</TypeIdentifier>
            </Grantee>
            <Permission>string</Permission>
          </S3Grant>
        </Grants>
        <Owner>
          <DisplayName>string</DisplayName>
          <ID>string</ID>
        </Owner>
      </AccessControlList>
      <CannedAccessControlList>string</CannedAccessControlList>
    </AccessControlPolicy>
  </S3PutObjectAcl>
  <S3PutObjectCopy>
    <AccessControlGrants>
      <S3Grant>
        <Grantee>
          <DisplayName>string</DisplayName>
          <Identifier>string</Identifier>
          <TypeIdentifier>string</TypeIdentifier>
        </Grantee>
        <Permission>string</Permission>
      </S3Grant>
    </AccessControlGrants>
    <AnnotationDirective>string</AnnotationDirective>
    <BucketKeyEnabled>boolean</BucketKeyEnabled>
    <CannedAccessControlList>string</CannedAccessControlList>
    <ChecksumAlgorithm>string</ChecksumAlgorithm>
    <MetadataDirective>string</MetadataDirective>
    <ModifiedSinceConstraint>timestamp</ModifiedSinceConstraint>
    <NewObjectMetadata>
      <CacheControl>string</CacheControl>
      <ContentDisposition>string</ContentDisposition>
      <ContentEncoding>string</ContentEncoding>
      <ContentLanguage>string</ContentLanguage>
      <ContentLength>long</ContentLength>
      <ContentMD5>string</ContentMD5>
    </NewObjectMetadata>
  </S3PutObjectCopy>

```

```

    <ContentType>string</ContentType>
    <HttpExpiresDate>timestamp</HttpExpiresDate>
    <RequesterCharged>boolean</RequesterCharged>
    <SSEAlgorithm>string</SSEAlgorithm>
    <UserMetadata>
        <entry>
            <key>string</key>
            <value>string</value>
        </entry>
    </UserMetadata>
</NewObjectMetadata>
<NewObjectTagging>
    <S3Tag>
        <Key>string</Key>
        <Value>string</Value>
    </S3Tag>
</NewObjectTagging>
<ObjectLockEventHold>string</ObjectLockEventHold>
<ObjectLockEventHoldDuration>
    <Days>integer</Days>
    <Years>integer</Years>
</ObjectLockEventHoldDuration>
<ObjectLockLegalHoldStatus>string</ObjectLockLegalHoldStatus>
<ObjectLockMode>string</ObjectLockMode>
<ObjectLockRetainUntilDate>timestamp</ObjectLockRetainUntilDate>
<RedirectLocation>string</RedirectLocation>
<RequesterPays>boolean</RequesterPays>
<SSEAwsKmsKeyId>string</SSEAwsKmsKeyId>
<StorageClass>string</StorageClass>
<TargetKeyPrefix>string</TargetKeyPrefix>
<TargetResource>string</TargetResource>
<UnmodifiedSinceConstraint>timestamp</UnmodifiedSinceConstraint>
</S3PutObjectCopy>
<S3PutObjectLegalHold>
    <LegalHold>
        <Status>string</Status>
    </LegalHold>
</S3PutObjectLegalHold>
<S3PutObjectRetention>
    <BypassGovernanceRetention>boolean</BypassGovernanceRetention>
    <Retention>
        <EventHold>string</EventHold>
        <EventHoldDuration>
            <Days>integer</Days>

```

```

        <Years>integer</Years>
    </EventHoldDuration>
    <Mode>string</Mode>
    <RetainUntilDate>timestamp</RetainUntilDate>
</Retention>
</S3PutObjectRetention>
<S3PutObjectTagging>
    <TagSet>
        <S3Tag>
            <Key>string</Key>
            <Value>string</Value>
        </S3Tag>
    </TagSet>
</S3PutObjectTagging>
<S3ReplicateObject>
</S3ReplicateObject>
<S3UpdateObjectEncryption>
    <ObjectEncryption>
        <SSE-KMS>
            <BucketKeyEnabled>boolean</BucketKeyEnabled>
            <KMSKeyArn>string</KMSKeyArn>
        </SSE-KMS>
    </ObjectEncryption>
</S3UpdateObjectEncryption>
</Operation>
<Priority>integer</Priority>
<ProgressSummary>
    <NumberOfTasksFailed>long</NumberOfTasksFailed>
    <NumberOfTasksSucceeded>long</NumberOfTasksSucceeded>
    <Timers>
        <ElapsedTimeInActiveSeconds>long</ElapsedTimeInActiveSeconds>
    </Timers>
    <TotalNumberOfTasks>long</TotalNumberOfTasks>
</ProgressSummary>
<Report>
    <Bucket>string</Bucket>
    <Enabled>boolean</Enabled>
    <ExpectedBucketOwner>string</ExpectedBucketOwner>
    <Format>string</Format>
    <Prefix>string</Prefix>
    <ReportScope>string</ReportScope>
</Report>
<RoleArn>string</RoleArn>
<Status>string</Status>

```

```
<StatusUpdateReason>string</StatusUpdateReason>
<SuspendedCause>string</SuspendedCause>
<SuspendedDate>timestamp</SuspendedDate>
<TerminationDate>timestamp</TerminationDate>
</Job>
</DescribeJobResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[DescribeJobResult](#)

Root level tag for the DescribeJobResult parameters.

Required: Yes

[Job](#)

Contains the configuration parameters and status for the job specified in the Describe Job request.

Type: [JobDescriptor](#) data type

Errors

BadRequestException

HTTP Status Code: 400

InternalServerErrorException

HTTP Status Code: 500

NotFoundException

HTTP Status Code: 400

TooManyRequestsException

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DescribeMultiRegionAccessPointOperation

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Retrieves the status of an asynchronous request to manage a Multi-Region Access Point. For more information about managing Multi-Region Access Points and how asynchronous requests work, see [Using Multi-Region Access Points](#) in the *Amazon S3 User Guide*.

The following actions are related to `GetMultiRegionAccessPoint`:

- [CreateMultiRegionAccessPoint](#)
- [DeleteMultiRegionAccessPoint](#)
- [GetMultiRegionAccessPoint](#)
- [ListMultiRegionAccessPoints](#)

Request Syntax

```
GET /v20180820/async-requests/mrap/request_token+ HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

[request_token](#)

The request token associated with the request you want to know about. This request token is returned as part of the response when you make an asynchronous request. You provide this token to query about the status of the asynchronous action.

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `arn: . +`

Required: Yes

[x-amz-account-id](#)

The AWS account ID for the owner of the Multi-Region Access Point.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<DescribeMultiRegionAccessPointOperationResult>
  <AsyncOperation>
    <CreationTime>timestamp</CreationTime>
    <Operation>string</Operation>
    <RequestParameters>
      <CreateMultiRegionAccessPointRequest>
        <Name>string</Name>
        <PublicAccessBlock>
          <BlockPublicAcls>boolean</BlockPublicAcls>
          <BlockPublicPolicy>boolean</BlockPublicPolicy>
          <IgnorePublicAcls>boolean</IgnorePublicAcls>
          <RestrictPublicBuckets>boolean</RestrictPublicBuckets>
        </PublicAccessBlock>
      </CreateMultiRegionAccessPointRequest>
      <Regions>
        <Region>
          <Bucket>string</Bucket>
          <BucketAccountId>string</BucketAccountId>
        </Region>
      </Regions>
    </DeleteMultiRegionAccessPointRequest>
    <Name>string</Name>
  </DeleteMultiRegionAccessPointRequest>
</DescribeMultiRegionAccessPointOperationResult>
```

```

<PutMultiRegionAccessPointPolicyRequest>
  <Name>string</Name>
  <Policy>string</Policy>
</PutMultiRegionAccessPointPolicyRequest>
</RequestParameters>
<RequestStatus>string</RequestStatus>
<RequestTokenARN>string</RequestTokenARN>
<ResponseDetails>
  <ErrorDetails>
    <Code>string</Code>
    <Message>string</Message>
    <RequestId>string</RequestId>
    <Resource>string</Resource>
  </ErrorDetails>
  <MultiRegionAccessPointDetails>
    <Regions>
      <Region>
        <Name>string</Name>
        <RequestStatus>string</RequestStatus>
      </Region>
    </Regions>
  </MultiRegionAccessPointDetails>
</ResponseDetails>
</AsyncOperation>
</DescribeMultiRegionAccessPointOperationResult>

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

DescribeMultiRegionAccessPointOperationResult

Root level tag for the DescribeMultiRegionAccessPointOperationResult parameters.

Required: Yes

AsyncOperation

A container element containing the details of the asynchronous operation.

Type: AsyncOperation data type

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DissociateAccessGrantsIdentityCenter

Service: Amazon S3 Control

Dissociates the AWS IAM Identity Center instance from the S3 Access Grants instance.

Permissions

You must have the `s3:DissociateAccessGrantsIdentityCenter` permission to use this operation.

Additional Permissions

You must have the `sso>DeleteApplication` permission to use this operation.

Request Syntax

```
DELETE /v20180820/accessgrantsinstance/identitycenter HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

[x-amz-account-id](#)

The AWS account ID of the S3 Access Grants instance.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetAccessGrant

Service: Amazon S3 Control

Get the details of an access grant from your S3 Access Grants instance.

Permissions

You must have the `s3:GetAccessGrant` permission to use this operation.

Request Syntax

```
GET /v20180820/accessgrantsinstance/grant/id HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

id

The ID of the access grant. S3 Access Grants auto-generates this ID when you create the access grant.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\ ]+`

Required: Yes

x-amz-account-id

The AWS account ID of the S3 Access Grants instance.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<GetAccessGrantResult>
  <CreatedAt>timestamp</CreatedAt>
  <AccessGrantId>string</AccessGrantId>
  <AccessGrantArn>string</AccessGrantArn>
  <Grantee>
    <GranteeIdentifier>string</GranteeIdentifier>
    <GranteeType>string</GranteeType>
  </Grantee>
  <Permission>string</Permission>
  <AccessGrantsLocationId>string</AccessGrantsLocationId>
  <AccessGrantsLocationConfiguration>
    <S3SubPrefix>string</S3SubPrefix>
  </AccessGrantsLocationConfiguration>
  <GrantScope>string</GrantScope>
  <ApplicationArn>string</ApplicationArn>
</GetAccessGrantResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

GetAccessGrantResult

Root level tag for the GetAccessGrantResult parameters.

Required: Yes

AccessGrantArn

The Amazon Resource Name (ARN) of the access grant.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:[a-z\-\-]+:s3:[a-z0-9\-\-]+:\d{12}:access\-\grants\grant/[a-zA-Z0-9\-\-]+`

AccessGrantId

The ID of the access grant. S3 Access Grants auto-generates this ID when you create the access grant.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [a-zA-Z0-9\-]+

AccessGrantsLocationConfiguration

The configuration options of the grant location. The grant location is the S3 path to the data to which you are granting access.

Type: [AccessGrantsLocationConfiguration](#) data type

AccessGrantsLocationId

The ID of the registered location to which you are granting access. S3 Access Grants assigns this ID when you register the location. S3 Access Grants assigns the ID default to the default location `s3://` and assigns an auto-generated ID to other locations that you register.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [a-zA-Z0-9\-]+

ApplicationArn

The Amazon Resource Name (ARN) of an AWS IAM Identity Center application associated with your Identity Center instance. If the grant includes an application ARN, the grantee can only access the S3 data through this application.

Type: String

Length Constraints: Minimum length of 10. Maximum length of 1224.

Pattern: arn:[^:]+:sso:\d{12}:application/.*\$

CreatedAt

The date and time when you created the access grant.

Type: Timestamp

Grantee

The user, group, or role to which you are granting access. You can grant access to an IAM user or role. If you have added a corporate directory to AWS IAM Identity Center and associated this Identity Center instance with the S3 Access Grants instance, the grantee can also be a corporate directory user or group.

Type: [Grantee](#) data type

GrantScope

The S3 path of the data to which you are granting access. It is the result of appending the `Subprefix` to the location scope.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2000.

Pattern: `^\.+$`

Permission

The type of permission that was granted in the access grant. Can be one of the following values:

- `READ` – Grant read-only access to the S3 data.
- `WRITE` – Grant write-only access to the S3 data.
- `READWRITE` – Grant both read and write access to the S3 data.

Type: String

Valid Values: `READ` | `WRITE` | `READWRITE`

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)

- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetAccessGrantsInstance

Service: Amazon S3 Control

Retrieves the S3 Access Grants instance for a Region in your account.

Permissions

You must have the `s3:GetAccessGrantsInstance` permission to use this operation.

Note

`GetAccessGrantsInstance` is not supported for cross-account access. You can only call the API from the account that owns the S3 Access Grants instance.

Request Syntax

```
GET /v20180820/accessgrantsinstance HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

[x-amz-account-id](#)

The AWS account ID of the S3 Access Grants instance.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<GetAccessGrantsInstanceResult>
  <AccessGrantsInstanceArn>string</AccessGrantsInstanceArn>
  <AccessGrantsInstanceId>string</AccessGrantsInstanceId>
  <IdentityCenterArn>string</IdentityCenterArn>
  <IdentityCenterInstanceArn>string</IdentityCenterInstanceArn>
  <IdentityCenterApplicationArn>string</IdentityCenterApplicationArn>
  <CreatedAt>timestamp</CreatedAt>
</GetAccessGrantsInstanceResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

GetAccessGrantsInstanceResult

Root level tag for the GetAccessGrantsInstanceResult parameters.

Required: Yes

AccessGrantsInstanceArn

The Amazon Resource Name (ARN) of the S3 Access Grants instance.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:[a-z\-\-]+:s3:[a-z0-9\-\-]+:\d{12}:access\-\grants\[a-zA-Z0-9\-\-]+`

AccessGrantsInstanceld

The ID of the S3 Access Grants instance. The ID is default. You can have one S3 Access Grants instance per Region per account.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\_]+`

CreatedAt

The date and time when you created the S3 Access Grants instance.

Type: Timestamp

IdentityCenterApplicationArn

If you associated your S3 Access Grants instance with an AWS IAM Identity Center instance, this field returns the Amazon Resource Name (ARN) of the IAM Identity Center instance application; a subresource of the original Identity Center instance. S3 Access Grants creates this Identity Center application for the specific S3 Access Grants instance.

Type: String

Length Constraints: Minimum length of 10. Maximum length of 1224.

Pattern: `arn:[^:]+\:sso::\d{12}:application/.*$`

IdentityCenterArn

This parameter has been deprecated.

If you associated your S3 Access Grants instance with an AWS IAM Identity Center instance, this field returns the Amazon Resource Name (ARN) of the IAM Identity Center instance application; a subresource of the original Identity Center instance. S3 Access Grants creates this Identity Center application for the specific S3 Access Grants instance.

Type: String

Length Constraints: Minimum length of 10. Maximum length of 1224.

Pattern: `arn:[^:]+\:sso::(\d{12}){0,1}:instance/.*$`

IdentityCenterInstanceArn

The Amazon Resource Name (ARN) of the AWS IAM Identity Center instance that you are associating with your S3 Access Grants instance. An IAM Identity Center instance is your corporate identity directory that you added to the IAM Identity Center. You can use the [ListInstances](#) API operation to retrieve a list of your Identity Center instances and their ARNs.

Type: String

Length Constraints: Minimum length of 10. Maximum length of 1224.

Pattern: `arn:[^:]+:sso:(\d{12}){0,1}:instance/.*$`

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetAccessGrantsInstanceForPrefix

Service: Amazon S3 Control

Retrieve the S3 Access Grants instance that contains a particular prefix.

Permissions

You must have the `s3:GetAccessGrantsInstanceForPrefix` permission for the caller account to use this operation.

Additional Permissions

The prefix owner account must grant you the following permissions to their S3 Access Grants instance: `s3:GetAccessGrantsInstanceForPrefix`.

Request Syntax

```
GET /v20180820/accessgrantsinstance/prefix?s3prefix=S3Prefix HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

s3prefix

The S3 prefix of the access grants that you would like to retrieve.

Length Constraints: Minimum length of 1. Maximum length of 2000.

Pattern: `^\.+`

Required: Yes

x-amz-account-id

The ID of the AWS account that is making this request.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<GetAccessGrantsInstanceForPrefixResult>
  <AccessGrantsInstanceArn>string</AccessGrantsInstanceArn>
  <AccessGrantsInstanceId>string</AccessGrantsInstanceId>
</GetAccessGrantsInstanceForPrefixResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

GetAccessGrantsInstanceForPrefixResult

Root level tag for the GetAccessGrantsInstanceForPrefixResult parameters.

Required: Yes

AccessGrantsInstanceArn

The Amazon Resource Name (ARN) of the S3 Access Grants instance.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:[a-z\-\-]+:s3:[a-z0-9\-\-]+:\d{12}:access\-\grants\[a-zA-Z0-9\-\-]+`

AccessGrantsInstanceId

The ID of the S3 Access Grants instance. The ID is default. You can have one S3 Access Grants instance per Region per account.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\ ]+`

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetAccessGrantsInstanceResourcePolicy

Service: Amazon S3 Control

Returns the resource policy of the S3 Access Grants instance.

Permissions

You must have the `s3:GetAccessGrantsInstanceResourcePolicy` permission to use this operation.

Request Syntax

```
GET /v20180820/accessgrantsinstance/resourcepolicy HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

x-amz-account-id

The AWS account ID of the S3 Access Grants instance.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<GetAccessGrantsInstanceResourcePolicyResult>
  <Policy>string</Policy>
```

```
<Organization>string</Organization>  
<CreatedAt>timestamp</CreatedAt>  
</GetAccessGrantsInstanceResourcePolicyResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[GetAccessGrantsInstanceResourcePolicyResult](#)

Root level tag for the GetAccessGrantsInstanceResourcePolicyResult parameters.

Required: Yes

[CreatedAt](#)

The date and time when you created the S3 Access Grants instance resource policy.

Type: Timestamp

[Organization](#)

The Organization of the resource policy of the S3 Access Grants instance.

Type: String

Length Constraints: Minimum length of 12. Maximum length of 34.

Pattern: ^o-[a-z0-9]{10,32}\$

[Policy](#)

The resource policy of the S3 Access Grants instance.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 350000.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetAccessGrantsLocation

Service: Amazon S3 Control

Retrieves the details of a particular location registered in your S3 Access Grants instance.

Permissions

You must have the `s3:GetAccessGrantsLocation` permission to use this operation.

Request Syntax

```
GET /v20180820/accessgrantsinstance/location/id HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

id

The ID of the registered location that you are retrieving. S3 Access Grants assigns this ID when you register the location. S3 Access Grants assigns the ID default to the default location `s3://` and assigns an auto-generated ID to other locations that you register.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\_]+`

Required: Yes

x-amz-account-id

The AWS account ID of the S3 Access Grants instance.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<GetAccessGrantsLocationResult>
  <CreatedAt>timestamp</CreatedAt>
  <AccessGrantsLocationId>string</AccessGrantsLocationId>
  <AccessGrantsLocationArn>string</AccessGrantsLocationArn>
  <LocationScope>string</LocationScope>
  <IAMRoleArn>string</IAMRoleArn>
</GetAccessGrantsLocationResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

GetAccessGrantsLocationResult

Root level tag for the GetAccessGrantsLocationResult parameters.

Required: Yes

AccessGrantsLocationArn

The Amazon Resource Name (ARN) of the registered location.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:[a-z\-\-]+:s3:[a-z0-9\-\-]+:\d{12}:access\-\grants\/location/[a-zA-Z0-9\-\-]+`

AccessGrantsLocationId

The ID of the registered location to which you are granting access. S3 Access Grants assigns this ID when you register the location. S3 Access Grants assigns the ID default to the default location `s3://` and assigns an auto-generated ID to other locations that you register.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\_]+`

CreatedAt

The date and time when you registered the location.

Type: Timestamp

IAMRoleArn

The Amazon Resource Name (ARN) of the IAM role for the registered location. S3 Access Grants assumes this role to manage access to the registered location.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:[^:]+\:iam:\d{12}:role/.*`

LocationScope

The S3 URI path to the registered location. The location scope can be the default S3 location `s3://`, the S3 path to a bucket, or the S3 path to a bucket and prefix. A prefix in S3 is a string of characters at the beginning of an object key name used to organize the objects that you store in your S3 buckets. For example, object key names that start with the `engineering/` prefix or object key names that start with the `marketing/campaigns/` prefix.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2000.

Pattern: `^\.+`

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)

- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetAccessPoint

Service: Amazon S3 Control

Returns configuration information about the specified access point.

All Amazon S3 on Outposts REST API requests for this action require an additional parameter of `x-amz-outpost-id` to be passed with the request. In addition, you must use an S3 on Outposts endpoint hostname prefix instead of `s3-control`. For an example of the request syntax for Amazon S3 on Outposts that uses the S3 on Outposts endpoint hostname prefix and the `x-amz-outpost-id` derived by using the access point ARN, see the [Examples](#) section.

The following actions are related to `GetAccessPoint`:

- [CreateAccessPoint](#)
- [DeleteAccessPoint](#)
- [ListAccessPoints](#)

Request Syntax

```
GET /v20180820/accesspoint/name HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the access point whose configuration information you want to retrieve.

For using this parameter with Amazon S3 on Outposts with the REST API, you must specify the name and the `x-amz-outpost-id` as well.

For using this parameter with S3 on Outposts with the AWS SDK and CLI, you must specify the ARN of the access point accessed in the format `arn:aws:s3-outposts:<Region>:<account-id>:outpost/<outpost-id>/accesspoint/<my-accesspoint-name>`. For example, to access the access point `reports-ap` through

Outpost my-outpost owned by account 123456789012 in Region us-west-2, use the URL encoding of `arn:aws:s3-outposts:us-west-2:123456789012:outpost/my-outpost/accesspoint/reports-ap`. The value must be URL encoded.

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

x-amz-account-id

The AWS account ID for the account that owns the specified access point.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<GetAccessPointResult>
  <Name>string</Name>
  <Bucket>string</Bucket>
  <NetworkOrigin>string</NetworkOrigin>
  <VpcConfiguration>
    <VpcId>string</VpcId>
  </VpcConfiguration>
  <PublicAccessBlockConfiguration>
    <BlockPublicAcls>boolean</BlockPublicAcls>
    <BlockPublicPolicy>boolean</BlockPublicPolicy>
    <IgnorePublicAcls>boolean</IgnorePublicAcls>
    <RestrictPublicBuckets>boolean</RestrictPublicBuckets>
  </PublicAccessBlockConfiguration>
  <CreationDate>timestamp</CreationDate>
  <Alias>string</Alias>
  <AccessPointArn>string</AccessPointArn>
  <Endpoints>
```

```
<entry>
  <key>string</key>
  <value>string</value>
</entry>
</Endpoints>
<BucketAccountId>string</BucketAccountId>
<DataSourceId>string</DataSourceId>
<DataSourceType>string</DataSourceType>
</GetAccessPointResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

GetAccessPointResult

Root level tag for the GetAccessPointResult parameters.

Required: Yes

AccessPointArn

The ARN of the access point.

Type: String

Length Constraints: Minimum length of 4. Maximum length of 128.

Alias

The name or alias of the access point.

Type: String

Length Constraints: Maximum length of 63.

Pattern: `^[0-9a-z\\-]{63}`

Bucket

The name of the bucket associated with the specified access point.

Type: String

Length Constraints: Maximum length of 255.

BucketAccountId

The AWS account ID associated with the S3 bucket associated with this access point.

Type: String

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

CreationDate

The date and time when the specified access point was created.

Type: Timestamp

DataSourceId

The unique identifier for the data source of the access point.

Type: String

Length Constraints: Maximum length of 191.

DataSourceType

The type of the data source that the access point is attached to.

Type: String

Endpoints

The VPC endpoint for the access point.

Type: String to string map

Key Length Constraints: Minimum length of 1. Maximum length of 64.

Value Length Constraints: Minimum length of 1. Maximum length of 1024.

Name

The name of the specified access point.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 255.

NetworkOrigin

Indicates whether this access point allows access from the public internet. If `VpcConfiguration` is specified for this access point, then `NetworkOrigin` is `VPC`, and the access point doesn't allow access from the public internet. Otherwise, `NetworkOrigin` is `Internet`, and the access point allows access from the public internet, subject to the access point and bucket access policies.

This will always be true for an Amazon S3 on Outposts access point

Type: String

Valid Values: `Internet` | `VPC`

PublicAccessBlockConfiguration

The `PublicAccessBlock` configuration that you want to apply to this Amazon S3 account. You can enable the configuration options in any combination. For more information about when Amazon S3 considers a bucket or object public, see [The Meaning of "Public"](#) in the *Amazon S3 User Guide*.

This data type is not supported for Amazon S3 on Outposts.

Type: [PublicAccessBlockConfiguration](#) data type

VpcConfiguration

Contains the virtual private cloud (VPC) configuration for the specified access point.

Note

This element is empty if this access point is an Amazon S3 on Outposts access point that is used by other AWS services.

Type: [VpcConfiguration](#) data type

Examples

Sample request syntax for getting an Amazon S3 on Outposts access point

The following request returns the access point of the specified S3 on Outposts access point.

```
GET /v20180820/accesspoint/example-access-point HTTP/1.1
Host: s3-outposts.<Region>.amazonaws.com
Date: Wed, 28 Oct 2020 22:32:00 GMT
Authorization: authorization string
x-amz-account-id: example-account-id
x-amz-outpost-id: op-01ac5d28a6a232904
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetAccessPointConfigurationForObjectLambda

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Returns configuration for an Object Lambda Access Point.

The following actions are related to `GetAccessPointConfigurationForObjectLambda`:

- [PutAccessPointConfigurationForObjectLambda](#)

Request Syntax

```
GET /v20180820/accesspointforobjectlambda/name/configuration HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the Object Lambda Access Point you want to return the configuration for.

Length Constraints: Minimum length of 3. Maximum length of 45.

Pattern: `^[a-z0-9]([a-z0-9\-*][a-z0-9])?$`

Required: Yes

x-amz-account-id

The account ID for the account that owns the specified Object Lambda Access Point.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<GetAccessPointConfigurationForObjectLambdaResult>
  <Configuration>
    <AllowedFeatures>
      <AllowedFeature>string</AllowedFeature>
    </AllowedFeatures>
    <CloudWatchMetricsEnabled>boolean</CloudWatchMetricsEnabled>
    <SupportingAccessPoint>string</SupportingAccessPoint>
    <TransformationConfigurations>
      <TransformationConfiguration>
        <Actions>
          <Action>string</Action>
        </Actions>
        <ContentTransformation>
          <AwsLambda>
            <FunctionArn>string</FunctionArn>
            <FunctionPayload>string</FunctionPayload>
          </AwsLambda>
        </ContentTransformation>
      </TransformationConfiguration>
    </TransformationConfigurations>
  </Configuration>
</GetAccessPointConfigurationForObjectLambdaResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

GetAccessPointConfigurationForObjectLambdaResult

Root level tag for the GetAccessPointConfigurationForObjectLambdaResult parameters.

Required: Yes

Configuration

Object Lambda Access Point configuration document.

Type: [ObjectLambdaConfiguration](#) data type

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetAccessPointForObjectLambda

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Returns configuration information about the specified Object Lambda Access Point

The following actions are related to `GetAccessPointForObjectLambda`:

- [CreateAccessPointForObjectLambda](#)
- [DeleteAccessPointForObjectLambda](#)
- [ListAccessPointsForObjectLambda](#)

Request Syntax

```
GET /v20180820/accesspointforobjectlambda/name HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

[name](#)

The name of the Object Lambda Access Point.

Length Constraints: Minimum length of 3. Maximum length of 45.

Pattern: `^[a-z0-9]([a-z0-9\-*][a-z0-9])?$`

Required: Yes

[x-amz-account-id](#)

The account ID for the account that owns the specified Object Lambda Access Point.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<GetAccessPointForObjectLambdaResult>
  <Name>string</Name>
  <PublicAccessBlockConfiguration>
    <BlockPublicAcls>boolean</BlockPublicAcls>
    <BlockPublicPolicy>boolean</BlockPublicPolicy>
    <IgnorePublicAcls>boolean</IgnorePublicAcls>
    <RestrictPublicBuckets>boolean</RestrictPublicBuckets>
  </PublicAccessBlockConfiguration>
  <CreationDate>timestamp</CreationDate>
  <Alias>
    <Status>string</Status>
    <Value>string</Value>
  </Alias>
</GetAccessPointForObjectLambdaResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

GetAccessPointForObjectLambdaResult

Root level tag for the GetAccessPointForObjectLambdaResult parameters.

Required: Yes

Alias

The alias of the Object Lambda Access Point.

Type: [ObjectLambdaAccessPointAlias](#) data type

CreationDate

The date and time when the specified Object Lambda Access Point was created.

Type: Timestamp

Name

The name of the Object Lambda Access Point.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 45.

Pattern: `^[a-z0-9]([a-z0-9\-*][a-z0-9])?$`

PublicAccessBlockConfiguration

Configuration to block all public access. This setting is turned on and can not be edited.

Type: [PublicAccessBlockConfiguration](#) data type

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetAccessPointPolicy

Service: Amazon S3 Control

Returns the access point policy associated with the specified access point.

The following actions are related to GetAccessPointPolicy:

- [PutAccessPointPolicy](#)
- [DeleteAccessPointPolicy](#)

Request Syntax

```
GET /v20180820/accesspoint/name/policy HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the access point whose policy you want to retrieve.

For using this parameter with Amazon S3 on Outposts with the REST API, you must specify the name and the x-amz-outpost-id as well.

For using this parameter with S3 on Outposts with the AWS SDK and CLI, you must specify the ARN of the access point accessed in the format `arn:aws:s3-outposts:<Region>:<account-id>:outpost/<outpost-id>/accesspoint/<my-accesspoint-name>`. For example, to access the access point `reports-ap` through Outpost `my-outpost` owned by account `123456789012` in Region `us-west-2`, use the URL encoding of `arn:aws:s3-outposts:us-west-2:123456789012:outpost/my-outpost/accesspoint/reports-ap`. The value must be URL encoded.

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

x-amz-account-id

The account ID for the account that owns the specified access point.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<GetAccessPointPolicyResult>
  <Policy>string</Policy>
</GetAccessPointPolicyResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

GetAccessPointPolicyResult

Root level tag for the GetAccessPointPolicyResult parameters.

Required: Yes

Policy

The access point policy associated with the specified access point.

Type: String

Examples

Sample request

The following request returns the access point of the specified Amazon S3 on Outposts.

```
GET /v20180820/accesspoint/example-access-point/policy HTTP/1.1
Host: s3-outposts.<Region>.amazonaws.com
Date: Wed, 28 Oct 2020 22:32:00 GMT
Authorization: authorization string
x-amz-account-id: 123456789012
x-amz-outpost-id: op-123456
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetAccessPointPolicyForObjectLambda

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Returns the resource policy for an Object Lambda Access Point.

The following actions are related to `GetAccessPointPolicyForObjectLambda`:

- [DeleteAccessPointPolicyForObjectLambda](#)
- [PutAccessPointPolicyForObjectLambda](#)

Request Syntax

```
GET /v20180820/accesspointforobjectlambda/name/policy HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the Object Lambda Access Point.

Length Constraints: Minimum length of 3. Maximum length of 45.

Pattern: `^[a-z0-9]([a-z0-9\-*][a-z0-9])?$`

Required: Yes

x-amz-account-id

The account ID for the account that owns the specified Object Lambda Access Point.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<GetAccessPointPolicyForObjectLambdaResult>
  <Policy>string</Policy>
</GetAccessPointPolicyForObjectLambdaResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

GetAccessPointPolicyForObjectLambdaResult

Root level tag for the GetAccessPointPolicyForObjectLambdaResult parameters.

Required: Yes

Policy

Object Lambda Access Point resource policy document.

Type: String

Examples

Sample resource policy

The following illustrates a sample resource policy.

```
{
```

```
"Version" : "2008-10-17",
"Statement": [{
  "Sid": "Grant account 123456789012 GetObject access",
  "Effect": "Allow",
  "Principal" : {
    "AWS": "arn:aws:iam::123456789012:root"
  },
  "Action": ["s3-object-lambda:GetObject"],
  "Resource": ["arn:aws:s3-object-lambda:us-east-1:123456789012:accesspoint/my-object-lambda-ap"]
},
{
  "Sid": "Grant account 444455556666 GetObject access",
  "Effect": "Allow",
  "Principal" : {
    "AWS": "arn:aws:iam::444455556666:root"
  },
  "Action": ["s3-object-lambda:GetObject"],
  "Resource": ["arn:aws:s3-object-lambda:us-east-1:123456789012:accesspoint/my-object-lambda-ap"]
}
]
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetAccessPointPolicyStatus

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Indicates whether the specified access point currently has a policy that allows public access. For more information about public access through access points, see [Managing Data Access with Amazon S3 access points](#) in the *Amazon S3 User Guide*.

Request Syntax

```
GET /v20180820/accesspoint/name/policyStatus HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the access point whose policy status you want to retrieve.

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

x-amz-account-id

The account ID for the account that owns the specified access point.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<GetAccessPointPolicyStatusResult>
  <PolicyStatus>
    <IsPublic>boolean</IsPublic>
  </PolicyStatus>
</GetAccessPointPolicyStatusResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

GetAccessPointPolicyStatusResult

Root level tag for the GetAccessPointPolicyStatusResult parameters.

Required: Yes

PolicyStatus

Indicates the current policy status of the specified access point.

Type: PolicyStatus data type

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetAccessPointPolicyStatusForObjectLambda

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Returns the status of the resource policy associated with an Object Lambda Access Point.

Request Syntax

```
GET /v20180820/accesspointforobjectlambda/name/policyStatus HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the Object Lambda Access Point.

Length Constraints: Minimum length of 3. Maximum length of 45.

Pattern: `^[a-z0-9]([a-z0-9\-*][a-z0-9])?$`

Required: Yes

x-amz-account-id

The account ID for the account that owns the specified Object Lambda Access Point.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<GetAccessPointPolicyStatusForObjectLambdaResult>
  <PolicyStatus>
    <IsPublic>boolean</IsPublic>
  </PolicyStatus>
</GetAccessPointPolicyStatusForObjectLambdaResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

GetAccessPointPolicyStatusForObjectLambdaResult

Root level tag for the GetAccessPointPolicyStatusForObjectLambdaResult parameters.

Required: Yes

PolicyStatus

Indicates whether this access point policy is public. For more information about how Amazon S3 evaluates policies to determine whether they are public, see [The Meaning of "Public"](#) in the *Amazon S3 User Guide*.

Type: [PolicyStatus](#) data type

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)

- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetAccessPointScope

Service: Amazon S3 Control

Returns the access point scope for a directory bucket.

To use this operation, you must have the permission to perform the `s3express:GetAccessPointScope` action.

For information about REST API errors, see [REST error responses](#).

Request Syntax

```
GET /v20180820/accesspoint/name/scope HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the access point with the scope you want to retrieve.

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

x-amz-account-id

The AWS account ID that owns the access point with the scope that you want to retrieve.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<GetAccessPointScopeResult>
  <Scope>
    <Permissions>
      <Permission>string</Permission>
    </Permissions>
    <Prefixes>
      <Prefix>string</Prefix>
    </Prefixes>
  </Scope>
</GetAccessPointScopeResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[GetAccessPointScopeResult](#)

Root level tag for the `GetAccessPointScopeResult` parameters.

Required: Yes

[Scope](#)

The contents of the access point scope.

Type: [Scope](#) data type

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)

- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucket

Service: Amazon S3 Control

Gets an Amazon S3 on Outposts bucket. For more information, see [Using Amazon S3 on Outposts](#) in the *Amazon S3 User Guide*.

If you are using an identity other than the root user of the AWS account that owns the Outposts bucket, the calling identity must have the `s3-outposts:GetBucket` permissions on the specified Outposts bucket and belong to the Outposts bucket owner's account in order to use this action. Only users from Outposts bucket owner account with the right permissions can perform actions on an Outposts bucket.

If you don't have `s3-outposts:GetBucket` permissions or you're not using an identity that belongs to the bucket owner's account, Amazon S3 returns a `403 Access Denied` error.

The following actions are related to `GetBucket` for Amazon S3 on Outposts:

All Amazon S3 on Outposts REST API requests for this action require an additional parameter of `x-amz-outpost-id` to be passed with the request. In addition, you must use an S3 on Outposts endpoint hostname prefix instead of `s3-control`. For an example of the request syntax for Amazon S3 on Outposts that uses the S3 on Outposts endpoint hostname prefix and the `x-amz-outpost-id` derived by using the access point ARN, see the [Examples](#) section.

- [PutObject](#)
- [CreateBucket](#)
- [DeleteBucket](#)

Request Syntax

```
GET /v20180820/bucket/name HTTP/1.1
Host: Bucket.s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

[name](#)

Specifies the bucket.

For using this parameter with Amazon S3 on Outposts with the REST API, you must specify the name and the `x-amz-outpost-id` as well.

For using this parameter with S3 on Outposts with the AWS SDK and CLI, you must specify the ARN of the bucket accessed in the format `arn:aws:s3-outposts:<Region>:<account-id>:outpost/<outpost-id>/bucket/<my-bucket-name>`. For example, to access the bucket `reports` through Outpost `my-outpost` owned by account `123456789012` in Region `us-west-2`, use the URL encoding of `arn:aws:s3-outposts:us-west-2:123456789012:outpost/my-outpost/bucket/reports`. The value must be URL encoded.

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

`x-amz-account-id`

The AWS account ID of the Outposts bucket.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<GetBucketResult>
  <Bucket>string</Bucket>
  <PublicAccessBlockEnabled>boolean</PublicAccessBlockEnabled>
  <CreationDate>timestamp</CreationDate>
</GetBucketResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

GetBucketResult

Root level tag for the GetBucketResult parameters.

Required: Yes

Bucket

The Outposts bucket requested.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 255.

CreationDate

The creation date of the Outposts bucket.

Type: Timestamp

PublicAccessBlockEnabled

Type: Boolean

Examples

Sample request for getting Amazon S3 on Outposts bucket

This example illustrates one usage of GetBucket.

```
GET /v20180820/bucket/example-outpost-bucket/ HTTP/1.1
Host: s3-outposts.<Region>.amazonaws.com
    x-amz-account-id: example-account-id
    x-amz-outpost-id: op-01ac5d28a6a232904
    x-amz-Date: 20200928T203757Z
    Authorization: authorization string
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketLifecycleConfiguration

Service: Amazon S3 Control

Note

This action gets an Amazon S3 on Outposts bucket's lifecycle configuration. To get an S3 bucket's lifecycle configuration, see [GetBucketLifecycleConfiguration](#) in the *Amazon S3 API Reference*.

Returns the lifecycle configuration information set on the Outposts bucket. For more information, see [Using Amazon S3 on Outposts](#) and for information about lifecycle configuration, see [Object Lifecycle Management](#) in *Amazon S3 User Guide*.

To use this action, you must have permission to perform the `s3-outposts:GetLifecycleConfiguration` action. The Outposts bucket owner has this permission, by default. The bucket owner can grant this permission to others. For more information about permissions, see [Permissions Related to Bucket Subresource Operations](#) and [Managing Access Permissions to Your Amazon S3 Resources](#).

All Amazon S3 on Outposts REST API requests for this action require an additional parameter of `x-amz-outpost-id` to be passed with the request. In addition, you must use an S3 on Outposts endpoint hostname prefix instead of `s3-control`. For an example of the request syntax for Amazon S3 on Outposts that uses the S3 on Outposts endpoint hostname prefix and the `x-amz-outpost-id` derived by using the access point ARN, see the [Examples](#) section.

`GetBucketLifecycleConfiguration` has the following special error:

- Error code: `NoSuchLifecycleConfiguration`
 - Description: The lifecycle configuration does not exist.
 - HTTP Status Code: 404 Not Found
 - SOAP Fault Code Prefix: Client

The following actions are related to `GetBucketLifecycleConfiguration`:

- [PutBucketLifecycleConfiguration](#)
- [DeleteBucketLifecycleConfiguration](#)

Request Syntax

```
GET /v20180820/bucket/name/lifecycleconfiguration HTTP/1.1
Host: Bucket.s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

name

The Amazon Resource Name (ARN) of the bucket.

For using this parameter with Amazon S3 on Outposts with the REST API, you must specify the name and the x-amz-outpost-id as well.

For using this parameter with S3 on Outposts with the AWS SDK and CLI, you must specify the ARN of the bucket accessed in the format `arn:aws:s3-outposts:<Region>:<account-id>:outpost/<outpost-id>/bucket/<my-bucket-name>`. For example, to access the bucket `reports` through Outpost `my-outpost` owned by account `123456789012` in Region `us-west-2`, use the URL encoding of `arn:aws:s3-outposts:us-west-2:123456789012:outpost/my-outpost/bucket/reports`. The value must be URL encoded.

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

x-amz-account-id

The AWS account ID of the Outposts bucket.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```

HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<GetBucketLifecycleConfigurationResult>
  <Rules>
    <Rule>
      <AbortIncompleteMultipartUpload>
        <DaysAfterInitiation>integer</DaysAfterInitiation>
      </AbortIncompleteMultipartUpload>
      <Expiration>
        <Date>timestamp</Date>
        <Days>integer</Days>
        <ExpiredObjectDeleteMarker>boolean</ExpiredObjectDeleteMarker>
      </Expiration>
      <Filter>
        <And>
          <ObjectSizeGreaterThan>long</ObjectSizeGreaterThan>
          <ObjectSizeLessThan>long</ObjectSizeLessThan>
          <Prefix>string</Prefix>
          <Tags>
            <S3Tag>
              <Key>string</Key>
              <Value>string</Value>
            </S3Tag>
          </Tags>
        </And>
        <ObjectSizeGreaterThan>long</ObjectSizeGreaterThan>
        <ObjectSizeLessThan>long</ObjectSizeLessThan>
        <Prefix>string</Prefix>
        <Tag>
          <Key>string</Key>
          <Value>string</Value>
        </Tag>
      </Filter>
      <ID>string</ID>
      <NoncurrentVersionExpiration>
        <NewerNoncurrentVersions>integer</NewerNoncurrentVersions>
        <NoncurrentDays>integer</NoncurrentDays>
      </NoncurrentVersionExpiration>
      <NoncurrentVersionTransitions>
        <NoncurrentVersionTransition>
          <NoncurrentDays>integer</NoncurrentDays>

```

```

        <StorageClass>string</StorageClass>
    </NoncurrentVersionTransition>
</NoncurrentVersionTransitions>
<Status>string</Status>
<Transitions>
    <Transition>
        <Date>timestamp</Date>
        <Days>integer</Days>
        <StorageClass>string</StorageClass>
    </Transition>
</Transitions>
</Rule>
</Rules>
</GetBucketLifecycleConfigurationResult>

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

GetBucketLifecycleConfigurationResult

Root level tag for the GetBucketLifecycleConfigurationResult parameters.

Required: Yes

Rules

Container for the lifecycle rule of the Outposts bucket.

Type: Array of [LifecycleRule](#) data types

Examples

Sample request to get the lifecycle configuration of the Amazon S3 on Outposts bucket

The following example shows how to get the lifecycle configuration of the Outposts bucket.

```

GET /v20180820/bucket/example-outpost-bucket/lifecycleconfiguration
HTTP/1.1
Host: s3-outposts.<Region>.amazonaws.com

```

```
x-amz-account-id: example-account-id
x-amz-outpost-id: op-01ac5d28a6a232904
x-amz-date: Thu, 15 Nov 2012 00:17:21 GMT
Authorization: signatureValue
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketPolicy

Service: Amazon S3 Control

Note

This action gets a bucket policy for an Amazon S3 on Outposts bucket. To get a policy for an S3 bucket, see [GetBucketPolicy](#) in the *Amazon S3 API Reference*.

Returns the policy of a specified Outposts bucket. For more information, see [Using Amazon S3 on Outposts](#) in the *Amazon S3 User Guide*.

If you are using an identity other than the root user of the AWS account that owns the bucket, the calling identity must have the `GetBucketPolicy` permissions on the specified bucket and belong to the bucket owner's account in order to use this action.

Only users from Outposts bucket owner account with the right permissions can perform actions on an Outposts bucket. If you don't have `s3-outposts:GetBucketPolicy` permissions or you're not using an identity that belongs to the bucket owner's account, Amazon S3 returns a 403 Access Denied error.

Important

As a security precaution, the root user of the AWS account that owns a bucket can always use this action, even if the policy explicitly denies the root user the ability to perform this action.

For more information about bucket policies, see [Using Bucket Policies and User Policies](#).

All Amazon S3 on Outposts REST API requests for this action require an additional parameter of `x-amz-outpost-id` to be passed with the request. In addition, you must use an S3 on Outposts endpoint hostname prefix instead of `s3-control`. For an example of the request syntax for Amazon S3 on Outposts that uses the S3 on Outposts endpoint hostname prefix and the `x-amz-outpost-id` derived by using the access point ARN, see the [Examples](#) section.

The following actions are related to `GetBucketPolicy`:

- [GetObject](#)

- [PutBucketPolicy](#)
- [DeleteBucketPolicy](#)

Request Syntax

```
GET /v20180820/bucket/name/policy HTTP/1.1
Host: Bucket.s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

[name](#)

Specifies the bucket.

For using this parameter with Amazon S3 on Outposts with the REST API, you must specify the name and the x-amz-outpost-id as well.

For using this parameter with S3 on Outposts with the AWS SDK and CLI, you must specify the ARN of the bucket accessed in the format `arn:aws:s3-outposts:<Region>:<account-id>:outpost/<outpost-id>/bucket/<my-bucket-name>`. For example, to access the bucket `reports` through Outpost `my-outpost` owned by account `123456789012` in Region `us-west-2`, use the URL encoding of `arn:aws:s3-outposts:us-west-2:123456789012:outpost/my-outpost/bucket/reports`. The value must be URL encoded.

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

[x-amz-account-id](#)

The AWS account ID of the Outposts bucket.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<GetBucketPolicyResult>
  <Policy>string</Policy>
</GetBucketPolicyResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[GetBucketPolicyResult](#)

Root level tag for the GetBucketPolicyResult parameters.

Required: Yes

[Policy](#)

The policy of the Outposts bucket.

Type: String

Examples

Sample GetBucketPolicy request for an Amazon S3 on Outposts bucket

The following request gets the policy of the specified Outposts bucket example-outpost-bucket.

```
GET /v20180820/bucket/example-outpost-bucket/policy HTTP/1.1
Host: s3-outposts.<Region>.amazonaws.com
Date: Wed, 28 Oct 2009 22:32:00 GMT
Authorization: authorization string
```

```
x-amz-account-id: example-account-id
x-amz-outpost-id: op-01ac5d28a6a232904
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketReplication

Service: Amazon S3 Control

Note

This operation gets an Amazon S3 on Outposts bucket's replication configuration. To get an S3 bucket's replication configuration, see [GetBucketReplication](#) in the *Amazon S3 API Reference*.

Returns the replication configuration of an S3 on Outposts bucket. For more information about S3 on Outposts, see [Using Amazon S3 on Outposts](#) in the *Amazon S3 User Guide*. For information about S3 replication on Outposts configuration, see [Replicating objects for S3 on Outposts](#) in the *Amazon S3 User Guide*.

Note

It can take a while to propagate PUT or DELETE requests for a replication configuration to all S3 on Outposts systems. Therefore, the replication configuration that's returned by a GET request soon after a PUT or DELETE request might return a more recent result than what's on the Outpost. If an Outpost is offline, the delay in updating the replication configuration on that Outpost can be significant.

This action requires permissions for the `s3-outposts:GetReplicationConfiguration` action. The Outposts bucket owner has this permission by default and can grant it to others. For more information about permissions, see [Setting up IAM with S3 on Outposts](#) and [Managing access to S3 on Outposts bucket](#) in the *Amazon S3 User Guide*.

All Amazon S3 on Outposts REST API requests for this action require an additional parameter of `x-amz-outpost-id` to be passed with the request. In addition, you must use an S3 on Outposts endpoint hostname prefix instead of `s3-control`. For an example of the request syntax for Amazon S3 on Outposts that uses the S3 on Outposts endpoint hostname prefix and the `x-amz-outpost-id` derived by using the access point ARN, see the [Examples](#) section.

If you include the `Filter` element in a replication configuration, you must also include the `DeleteMarkerReplication`, `Status`, and `Priority` elements. The response also returns those elements.

For information about S3 on Outposts replication failure reasons, see [Replication failure reasons](#) in the *Amazon S3 User Guide*.

The following operations are related to `GetBucketReplication`:

- [PutBucketReplication](#)
- [DeleteBucketReplication](#)

Request Syntax

```
GET /v20180820/bucket/name/replication HTTP/1.1
Host: Bucket.s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

[name](#)

Specifies the bucket to get the replication information for.

For using this parameter with Amazon S3 on Outposts with the REST API, you must specify the name and the `x-amz-outpost-id` as well.

For using this parameter with S3 on Outposts with the AWS SDK and CLI, you must specify the ARN of the bucket accessed in the format `arn:aws:s3-outposts:<Region>:<account-id>:outpost/<outpost-id>/bucket/<my-bucket-name>`. For example, to access the bucket reports through Outpost my-outpost owned by account 123456789012 in Region us-west-2, use the URL encoding of `arn:aws:s3-outposts:us-west-2:123456789012:outpost/my-outpost/bucket/reports`. The value must be URL encoded.

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

[x-amz-account-id](#)

The AWS account ID of the Outposts bucket.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<GetBucketReplicationResult>
  <ReplicationConfiguration>
    <Role>string</Role>
    <Rules>
      <Rule>
        <Bucket>string</Bucket>
        <DeleteMarkerReplication>
          <Status>string</Status>
        </DeleteMarkerReplication>
        <Destination>
          <AccessControlTranslation>
            <Owner>string</Owner>
          </AccessControlTranslation>
          <Account>string</Account>
          <Bucket>string</Bucket>
          <EncryptionConfiguration>
            <ReplicaKmsKeyID>string</ReplicaKmsKeyID>
          </EncryptionConfiguration>
          <Metrics>
            <EventThreshold>
              <Minutes>integer</Minutes>
            </EventThreshold>
            <Status>string</Status>
          </Metrics>
          <ReplicationTime>
            <Status>string</Status>
            <Time>
              <Minutes>integer</Minutes>
            </Time>
          </ReplicationTime>
        </Destination>
      </Rule>
    </Rules>
  </ReplicationConfiguration>
</GetBucketReplicationResult>
```

```

    </ReplicationTime>
    <StorageClass>string</StorageClass>
  </Destination>
  <ExistingObjectReplication>
    <Status>string</Status>
  </ExistingObjectReplication>
  <Filter>
    <And>
      <Prefix>string</Prefix>
      <Tags>
        <S3Tag>
          <Key>string</Key>
          <Value>string</Value>
        </S3Tag>
      </Tags>
    </And>
    <Prefix>string</Prefix>
    <Tag>
      <Key>string</Key>
      <Value>string</Value>
    </Tag>
  </Filter>
  <ID>string</ID>
  <Prefix>string</Prefix>
  <Priority>integer</Priority>
  <SourceSelectionCriteria>
    <ReplicaModifications>
      <Status>string</Status>
    </ReplicaModifications>
    <SseKmsEncryptedObjects>
      <Status>string</Status>
    </SseKmsEncryptedObjects>
  </SourceSelectionCriteria>
  <Status>string</Status>
</Rule>
</Rules>
</ReplicationConfiguration>
</GetBucketReplicationResult>

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

GetBucketReplicationResult

Root level tag for the GetBucketReplicationResult parameters.

Required: Yes

ReplicationConfiguration

A container for one or more replication rules. A replication configuration must have at least one rule and you can add up to 100 rules. The maximum size of a replication configuration is 128 KB.

Type: [ReplicationConfiguration](#) data type

Examples

Sample request to get the replication configuration of an Amazon S3 on Outposts bucket

The following example shows how to get the replication configuration of an Outposts bucket.

```
GET /v20180820/bucket/example-outpost-bucket/replication HTTP/1.1
Host: s3-outposts.<Region>.amazonaws.com
x-amz-account-id: example-account-id
x-amz-outpost-id: op-01ac5d28a6a232904
Authorization: signatureValue
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)

- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketTagging

Service: Amazon S3 Control

Note

This action gets an Amazon S3 on Outposts bucket's tags. To get an S3 bucket tags, see [GetBucketTagging](#) in the *Amazon S3 API Reference*.

Returns the tag set associated with the Outposts bucket. For more information, see [Using Amazon S3 on Outposts](#) in the *Amazon S3 User Guide*.

To use this action, you must have permission to perform the GetBucketTagging action. By default, the bucket owner has this permission and can grant this permission to others.

GetBucketTagging has the following special error:

- Error code: NoSuchTagSetError
 - Description: There is no tag set associated with the bucket.

All Amazon S3 on Outposts REST API requests for this action require an additional parameter of `x-amz-outpost-id` to be passed with the request. In addition, you must use an S3 on Outposts endpoint hostname prefix instead of `s3-control`. For an example of the request syntax for Amazon S3 on Outposts that uses the S3 on Outposts endpoint hostname prefix and the `x-amz-outpost-id` derived by using the access point ARN, see the [Examples](#) section.

The following actions are related to GetBucketTagging:

- [PutBucketTagging](#)
- [DeleteBucketTagging](#)

Request Syntax

```
GET /v20180820/bucket/name/tagging HTTP/1.1
Host: Bucket.s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

name

Specifies the bucket.

For using this parameter with Amazon S3 on Outposts with the REST API, you must specify the name and the x-amz-outpost-id as well.

For using this parameter with S3 on Outposts with the AWS SDK and CLI, you must specify the ARN of the bucket accessed in the format `arn:aws:s3-outposts:<Region>:<account-id>:outpost/<outpost-id>/bucket/<my-bucket-name>`. For example, to access the bucket `reports` through Outpost `my-outpost` owned by account `123456789012` in Region `us-west-2`, use the URL encoding of `arn:aws:s3-outposts:us-west-2:123456789012:outpost/my-outpost/bucket/reports`. The value must be URL encoded.

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

x-amz-account-id

The AWS account ID of the Outposts bucket.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
```

```
<GetBucketTaggingResult>
  <TagSet>
    <S3Tag>
      <Key>string</Key>
      <Value>string</Value>
    </S3Tag>
  </TagSet>
</GetBucketTaggingResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[GetBucketTaggingResult](#)

Root level tag for the GetBucketTaggingResult parameters.

Required: Yes

[TagSet](#)

The tags set of the Outposts bucket.

Type: Array of [S3Tag](#) data types

Examples

Amazon S3 on Outposts request example for getting a tag set for an Outposts bucket

The following request gets the tag set of the specified Outposts bucket example-outpost-bucket.

```
GET /v20180820/bucket/example-outpost-bucket/tagging HTTP/1.1
Host: s3-outposts.<Region>.amazonaws.com
x-amz-date: Wed, 28 Oct 2020 22:32:00 GMT
x-amz-account-id: example-account-id
x-amz-outpost-id: op-01ac5d28a6a232904
Authorization: authorization string
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetBucketVersioning

Service: Amazon S3 Control

Note

This operation returns the versioning state for S3 on Outposts buckets only. To return the versioning state for an S3 bucket, see [GetBucketVersioning](#) in the *Amazon S3 API Reference*.

Returns the versioning state for an S3 on Outposts bucket. With S3 Versioning, you can save multiple distinct copies of your objects and recover from unintended user actions and application failures.

If you've never set versioning on your bucket, it has no versioning state. In that case, the `GetBucketVersioning` request does not return a versioning state value.

For more information about versioning, see [Versioning](#) in the *Amazon S3 User Guide*.

All Amazon S3 on Outposts REST API requests for this action require an additional parameter of `x-amz-outpost-id` to be passed with the request. In addition, you must use an S3 on Outposts endpoint hostname prefix instead of `s3-control`. For an example of the request syntax for Amazon S3 on Outposts that uses the S3 on Outposts endpoint hostname prefix and the `x-amz-outpost-id` derived by using the access point ARN, see the [Examples](#) section.

The following operations are related to `GetBucketVersioning` for S3 on Outposts.

- [PutBucketVersioning](#)
- [PutBucketLifecycleConfiguration](#)
- [GetBucketLifecycleConfiguration](#)

Request Syntax

```
GET /v20180820/bucket/name/versioning HTTP/1.1
Host: Bucket.s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

name

The S3 on Outposts bucket to return the versioning state for.

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

x-amz-account-id

The AWS account ID of the S3 on Outposts bucket.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<GetBucketVersioningResult>
  <Status>string</Status>
  <MfaDelete>string</MfaDelete>
</GetBucketVersioningResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

GetBucketVersioningResult

Root level tag for the GetBucketVersioningResult parameters.

Required: Yes

MFADelete

Specifies whether MFA delete is enabled in the bucket versioning configuration. This element is returned only if the bucket has been configured with MFA delete. If MFA delete has never been configured for the bucket, this element is not returned.

Type: String

Valid Values: Enabled | Disabled

Status

The versioning state of the S3 on Outposts bucket.

Type: String

Valid Values: Enabled | Suspended

Examples

Sample GetBucketVersioning request on an S3 on Outposts bucket

This request returns the versioning state for an S3 on Outposts bucket that's named example-outpost-bucket.

```
GET /v20180820/bucket/example-outpost-bucket/?versioning HTTP/1.1
Host:s3-outposts.region-code.amazonaws.com
x-amz-account-id: example-account-id
x-amz-outpost-id: op-01ac5d28a6a232904
x-amz-date: Wed, 25 May 2022 00:14:21 GMT
Authorization: signatureValue
```

Sample GetBucketVersioning response on a versioning-enabled S3 on Outposts bucket

If you enabled versioning on a bucket, the response is:

```
<VersioningConfiguration xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <Status>Enabled</Status>
```

```
</VersioningConfiguration>
```

Sample GetBucketVersioning response on a versioning-suspended bucket

If you suspended versioning on a bucket, the response is:

```
<VersioningConfiguration xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <Status>Suspended</Status>
</VersioningConfiguration>
```

Sample GetBucketVersioning response if you have never enabled versioning.

If you have never enabled versioning on a bucket, the response is:

```
<VersioningConfiguration xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetDataAccess

Service: Amazon S3 Control

Returns a temporary access credential from S3 Access Grants to the grantee or client application. The [temporary credential](#) is an AWS STS token that grants them access to the S3 data.

Permissions

You must have the `s3:GetDataAccess` permission to use this operation.

Additional Permissions

The IAM role that S3 Access Grants assumes must have the following permissions specified in the trust policy when registering the location: `sts:AssumeRole`, for directory users or groups `sts:SetContext`, and for IAM users or roles `sts:SetSourceIdentity`.

Request Syntax

```
GET /v20180820/accessgrantsinstance/dataaccess?  
auditContext=AuditContext&durationSeconds=DurationSeconds&permission=Permission&privilege=Privilege  
HTTP/1.1  
Host: s3-control.amazonaws.com  
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

auditContext

The context to identify the job or query associated with the credential request. This information will be displayed in CloudTrail log in your account.

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `[\\u0020-\\uD7FF\\uE000-\\uFFFF\\uD800\\uDC00-\\uDBFF\\uDFFF\\t]*`

durationSeconds

The session duration, in seconds, of the temporary access credential that S3 Access Grants vends to the grantee or client application. The default value is 1 hour, but the grantee can

specify a range from 900 seconds (15 minutes) up to 43200 seconds (12 hours). If the grantee requests a value higher than this maximum, the operation fails.

Valid Range: Minimum value of 900. Maximum value of 43200.

[permission](#)

The type of permission granted to your S3 data, which can be set to one of the following values:

- **READ** – Grant read-only access to the S3 data.
- **WRITE** – Grant write-only access to the S3 data.
- **READWRITE** – Grant both read and write access to the S3 data.

Valid Values: **READ** | **WRITE** | **READWRITE**

Required: Yes

[privilege](#)

The scope of the temporary access credential that S3 Access Grants vends to the grantee or client application.

- **Default** – The scope of the returned temporary access token is the scope of the grant that is closest to the target scope.
- **Minimal** – The scope of the returned temporary access token is the same as the requested target scope as long as the requested scope is the same as or a subset of the grant scope.

Valid Values: **Minimal** | **Default**

[target](#)

The S3 URI path of the data to which you are requesting temporary access credentials. If the requesting account has an access grant for this data, S3 Access Grants vends temporary access credentials in the response.

Length Constraints: Minimum length of 1. Maximum length of 2000.

Pattern: `^.+`

Required: Yes

targetType

The type of Target. The only possible value is `Object`. Pass this value if the target data that you would like to access is a path to an object. Do not pass this value if the target data is a bucket or a bucket and a prefix.

Valid Values: `Object`

x-amz-account-id

The AWS account ID of the S3 Access Grants instance.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<GetDataAccessResult>
  <Credentials>
    <AccessKeyId>string</AccessKeyId>
    <Expiration>timestamp</Expiration>
    <SecretAccessKey>string</SecretAccessKey>
    <SessionToken>string</SessionToken>
  </Credentials>
  <MatchedGrantTarget>string</MatchedGrantTarget>
  <Grantee>
    <GranteeIdentifier>string</GranteeIdentifier>
    <GranteeType>string</GranteeType>
  </Grantee>
</GetDataAccessResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

GetDataAccessResult

Root level tag for the GetDataAccessResult parameters.

Required: Yes

Credentials

The temporary credential token that S3 Access Grants vends.

Type: [Credentials](#) data type

Grantee

The user, group, or role that was granted access to the S3 location scope. For directory identities, this API also returns the grants of the IAM role used for the identity-aware request. For more information on identity-aware sessions, see [Granting permissions to use identity-aware console sessions](#).

Type: [Grantee](#) data type

MatchedGrantTarget

The S3 URI path of the data to which you are being granted temporary access credentials.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2000.

Pattern: ^.+ \$

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetJobTagging

Service: Amazon S3 Control

Returns the tags on an S3 Batch Operations job.

Permissions

To use the GetJobTagging operation, you must have permission to perform the `s3:GetJobTagging` action. For more information, see [Controlling access and labeling jobs using tags](#) in the *Amazon S3 User Guide*.

Related actions include:

- [CreateJob](#)
- [PutJobTagging](#)
- [DeleteJobTagging](#)

Request Syntax

```
GET /v20180820/jobs/id/tagging HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

id

The ID for the S3 Batch Operations job whose tags you want to retrieve.

Length Constraints: Minimum length of 5. Maximum length of 36.

Pattern: `[a-zA-Z0-9\-\_\_]+`

Required: Yes

x-amz-account-id

The AWS account ID associated with the S3 Batch Operations job.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<GetJobTaggingResult>
  <Tags>
    <S3Tag>
      <Key>string</Key>
      <Value>string</Value>
    </S3Tag>
  </Tags>
</GetJobTaggingResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[GetJobTaggingResult](#)

Root level tag for the `GetJobTaggingResult` parameters.

Required: Yes

[Tags](#)

The set of tags associated with the S3 Batch Operations job.

Type: Array of [S3Tag](#) data types

Errors

InternalServiceException

HTTP Status Code: 500

NotFoundException

HTTP Status Code: 400

TooManyRequestsException

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetMultiRegionAccessPoint

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Returns configuration information about the specified Multi-Region Access Point.

This action will always be routed to the US West (Oregon) Region. For more information about the restrictions around working with Multi-Region Access Points, see [Multi-Region Access Point restrictions and limitations](#) in the *Amazon S3 User Guide*.

The following actions are related to GetMultiRegionAccessPoint:

- [CreateMultiRegionAccessPoint](#)
- [DeleteMultiRegionAccessPoint](#)
- [DescribeMultiRegionAccessPointOperation](#)
- [ListMultiRegionAccessPoints](#)

Request Syntax

```
GET /v20180820/mrap/instances/name+ HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the Multi-Region Access Point whose configuration information you want to receive. The name of the Multi-Region Access Point is different from the alias. For more information about the distinction between the name and the alias of an Multi-Region Access Point, see [Rules for naming Amazon S3 Multi-Region Access Points](#) in the *Amazon S3 User Guide*.

Length Constraints: Maximum length of 50.

Pattern: `^[a-z0-9][-a-z0-9]{1,48}[a-z0-9]$`

Required: Yes

[x-amz-account-id](#)

The AWS account ID for the owner of the Multi-Region Access Point.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<GetMultiRegionAccessPointResult>
  <AccessPoint>
    <Alias>string</Alias>
    <CreatedAt>timestamp</CreatedAt>
    <Name>string</Name>
    <PublicAccessBlock>
      <BlockPublicAcls>boolean</BlockPublicAcls>
      <BlockPublicPolicy>boolean</BlockPublicPolicy>
      <IgnorePublicAcls>boolean</IgnorePublicAcls>
      <RestrictPublicBuckets>boolean</RestrictPublicBuckets>
    </PublicAccessBlock>
    <Regions>
      <Region>
        <Bucket>string</Bucket>
        <BucketAccountId>string</BucketAccountId>
        <Region>string</Region>
      </Region>
    </Regions>
  <Status>string</Status>
```

```
</AccessPoint>  
</GetMultiRegionAccessPointResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

GetMultiRegionAccessPointResult

Root level tag for the GetMultiRegionAccessPointResult parameters.

Required: Yes

AccessPoint

A container element containing the details of the requested Multi-Region Access Point.

Type: [MultiRegionAccessPointReport](#) data type

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetMultiRegionAccessPointPolicy

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Returns the access control policy of the specified Multi-Region Access Point.

This action will always be routed to the US West (Oregon) Region. For more information about the restrictions around working with Multi-Region Access Points, see [Multi-Region Access Point restrictions and limitations](#) in the *Amazon S3 User Guide*.

The following actions are related to GetMultiRegionAccessPointPolicy:

- [GetMultiRegionAccessPointPolicyStatus](#)
- [PutMultiRegionAccessPointPolicy](#)

Request Syntax

```
GET /v20180820/mrap/instances/name+/policy HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

name

Specifies the Multi-Region Access Point. The name of the Multi-Region Access Point is different from the alias. For more information about the distinction between the name and the alias of an Multi-Region Access Point, see [Rules for naming Amazon S3 Multi-Region Access Points](#) in the *Amazon S3 User Guide*.

Length Constraints: Maximum length of 50.

Pattern: `^[a-z0-9][-a-z0-9]{1,48}[a-z0-9]$`

Required: Yes

[x-amz-account-id](#)

The AWS account ID for the owner of the Multi-Region Access Point.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<GetMultiRegionAccessPointPolicyResult>
  <Policy>
    <Established>
      <Policy>string</Policy>
    </Established>
    <Proposed>
      <Policy>string</Policy>
    </Proposed>
  </Policy>
</GetMultiRegionAccessPointPolicyResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[GetMultiRegionAccessPointPolicyResult](#)

Root level tag for the `GetMultiRegionAccessPointPolicyResult` parameters.

Required: Yes

Policy

The policy associated with the specified Multi-Region Access Point.

Type: [MultiRegionAccessPointPolicyDocument](#) data type

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetMultiRegionAccessPointPolicyStatus

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Indicates whether the specified Multi-Region Access Point has an access control policy that allows public access.

This action will always be routed to the US West (Oregon) Region. For more information about the restrictions around working with Multi-Region Access Points, see [Multi-Region Access Point restrictions and limitations](#) in the *Amazon S3 User Guide*.

The following actions are related to GetMultiRegionAccessPointPolicyStatus:

- [GetMultiRegionAccessPointPolicy](#)
- [PutMultiRegionAccessPointPolicy](#)

Request Syntax

```
GET /v20180820/mrap/instances/name+/policystatus HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

name

Specifies the Multi-Region Access Point. The name of the Multi-Region Access Point is different from the alias. For more information about the distinction between the name and the alias of an Multi-Region Access Point, see [Rules for naming Amazon S3 Multi-Region Access Points](#) in the *Amazon S3 User Guide*.

Length Constraints: Maximum length of 50.

Pattern: `^[a-z0-9][-a-z0-9]{1,48}[a-z0-9]$`

Required: Yes

[x-amz-account-id](#)

The AWS account ID for the owner of the Multi-Region Access Point.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<GetMultiRegionAccessPointPolicyStatusResult>
  <Established>
    <IsPublic>boolean</IsPublic>
  </Established>
</GetMultiRegionAccessPointPolicyStatusResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[GetMultiRegionAccessPointPolicyStatusResult](#)

Root level tag for the `GetMultiRegionAccessPointPolicyStatusResult` parameters.

Required: Yes

[Established](#)

Indicates whether this access point policy is public. For more information about how Amazon S3 evaluates policies to determine whether they are public, see [The Meaning of "Public"](#) in the *Amazon S3 User Guide*.

Type: [PolicyStatus](#) data type

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetMultiRegionAccessPointRoutes

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Returns the routing configuration for a Multi-Region Access Point, indicating which Regions are active or passive.

To obtain routing control changes and failover requests, use the Amazon S3 failover control infrastructure endpoints in these five AWS Regions:

- us-east-1
- us-west-2
- ap-southeast-2
- ap-northeast-1
- eu-west-1

Request Syntax

```
GET /v20180820/mrap/instances/mrp+/routes HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

mrp

The Multi-Region Access Point ARN.

Length Constraints: Maximum length of 200.

Pattern: `^[a-zA-Z0-9\:\.-]{3,200}$`

Required: Yes

[x-amz-account-id](#)

The AWS account ID for the owner of the Multi-Region Access Point.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<GetMultiRegionAccessPointRoutesResult>
  <Mrap>string</Mrap>
  <Routes>
    <Route>
      <Bucket>string</Bucket>
      <Region>string</Region>
      <TrafficDialPercentage>integer</TrafficDialPercentage>
    </Route>
  </Routes>
</GetMultiRegionAccessPointRoutesResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[GetMultiRegionAccessPointRoutesResult](#)

Root level tag for the `GetMultiRegionAccessPointRoutesResult` parameters.

Required: Yes

Mrap

The Multi-Region Access Point ARN.

Type: String

Length Constraints: Maximum length of 200.

Pattern: `^[a-zA-Z0-9\:\.-]{3,200}$`

Routes

The different routes that make up the route configuration. Active routes return a value of 100, and passive routes return a value of 0.

Type: Array of [MultiRegionAccessPointRoute](#) data types

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetPublicAccessBlock

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Retrieves the `PublicAccessBlock` configuration for an AWS account. This operation returns the effective account-level configuration, which may inherit from organization-level policies. For more information, see [Using Amazon S3 block public access](#).

Related actions include:

- [DeletePublicAccessBlock](#)
- [PutPublicAccessBlock](#)

Request Syntax

```
GET /v20180820/configuration/publicAccessBlock HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

[x-amz-account-id](#)

The account ID for the AWS account whose `PublicAccessBlock` configuration you want to retrieve.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<PublicAccessBlockConfiguration>
  <BlockPublicAcls>boolean</BlockPublicAcls>
  <IgnorePublicAcls>boolean</IgnorePublicAcls>
  <BlockPublicPolicy>boolean</BlockPublicPolicy>
  <RestrictPublicBuckets>boolean</RestrictPublicBuckets>
</PublicAccessBlockConfiguration>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

PublicAccessBlockConfiguration

Root level tag for the PublicAccessBlockConfiguration parameters.

Required: Yes

BlockPublicAcls

Specifies whether Amazon S3 should block public access control lists (ACLs) for buckets in this account. Setting this element to TRUE causes the following behavior:

- PutBucketAcl and PutObjectAcl calls fail if the specified ACL is public.
- PUT Object calls fail if the request includes a public ACL.
- PUT Bucket calls fail if the request includes a public ACL.

Enabling this setting doesn't affect existing policies or ACLs.

This property is not supported for Amazon S3 on Outposts.

Type: Boolean

BlockPublicPolicy

Specifies whether Amazon S3 should block public bucket policies for buckets in this account. Setting this element to TRUE causes Amazon S3 to reject calls to PUT Bucket policy if the specified bucket policy allows public access.

Enabling this setting doesn't affect existing bucket policies.

This property is not supported for Amazon S3 on Outposts.

Type: Boolean

IgnorePublicAcls

Specifies whether Amazon S3 should ignore public ACLs for buckets in this account. Setting this element to TRUE causes Amazon S3 to ignore all public ACLs on buckets in this account and any objects that they contain.

Enabling this setting doesn't affect the persistence of any existing ACLs and doesn't prevent new public ACLs from being set.

This property is not supported for Amazon S3 on Outposts.

Type: Boolean

RestrictPublicBuckets

Specifies whether Amazon S3 should restrict public bucket policies for buckets in this account. Setting this element to TRUE restricts access to buckets with public policies to only AWS service principals and authorized users within this account.

Enabling this setting doesn't affect previously stored bucket policies, except that public and cross-account access within any public bucket policy, including non-public delegation to specific accounts, is blocked.

This property is not supported for Amazon S3 on Outposts.

Type: Boolean

Errors

NoSuchPublicAccessBlockConfiguration

Amazon S3 throws this exception if you make a `GetPublicAccessBlock` request against an account that doesn't have a `PublicAccessBlockConfiguration` set.

HTTP Status Code: 404

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetStorageLensConfiguration

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Gets the Amazon S3 Storage Lens configuration. For more information, see [Assessing your storage activity and usage with Amazon S3 Storage Lens](#) in the *Amazon S3 User Guide*. For a complete list of S3 Storage Lens metrics, see [S3 Storage Lens metrics glossary](#) in the *Amazon S3 User Guide*.

Note

To use this action, you must have permission to perform the `s3:GetStorageLensConfiguration` action. For more information, see [Setting permissions to use Amazon S3 Storage Lens](#) in the *Amazon S3 User Guide*.

Request Syntax

```
GET /v20180820/storageLens/storagelensid HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

storagelensid

The ID of the Amazon S3 Storage Lens configuration.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\_\.\ ]+`

Required: Yes

x-amz-account-id

The account ID of the requester.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<StorageLensConfiguration>
  <Id>string</Id>
  <AccountLevel>
    <ActivityMetrics>
      <IsEnabled>boolean</IsEnabled>
    </ActivityMetrics>
    <AdvancedCostOptimizationMetrics>
      <IsEnabled>boolean</IsEnabled>
    </AdvancedCostOptimizationMetrics>
    <AdvancedDataProtectionMetrics>
      <IsEnabled>boolean</IsEnabled>
    </AdvancedDataProtectionMetrics>
    <AdvancedPerformanceMetrics>
      <IsEnabled>boolean</IsEnabled>
    </AdvancedPerformanceMetrics>
  <BucketLevel>
    <ActivityMetrics>
      <IsEnabled>boolean</IsEnabled>
    </ActivityMetrics>
    <AdvancedCostOptimizationMetrics>
      <IsEnabled>boolean</IsEnabled>
    </AdvancedCostOptimizationMetrics>
    <AdvancedDataProtectionMetrics>
      <IsEnabled>boolean</IsEnabled>
    </AdvancedDataProtectionMetrics>
    <AdvancedPerformanceMetrics>
      <IsEnabled>boolean</IsEnabled>
    </AdvancedPerformanceMetrics>
    <DetailedStatusCodesMetrics>
```

```

    <IsEnabled>boolean</IsEnabled>
  </DetailedStatusCodesMetrics>
  <PrefixLevel>
    <StorageMetrics>
      <IsEnabled>boolean</IsEnabled>
      <SelectionCriteria>
        <Delimiter>string</Delimiter>
        <MaxDepth>integer</MaxDepth>
        <MinStorageBytesPercentage>double</MinStorageBytesPercentage>
      </SelectionCriteria>
    </StorageMetrics>
  </PrefixLevel>
</BucketLevel>
<DetailedStatusCodesMetrics>
  <IsEnabled>boolean</IsEnabled>
</DetailedStatusCodesMetrics>
<StorageLensGroupLevel>
  <SelectionCriteria>
    <Exclude>
      <Arn>string</Arn>
    </Exclude>
    <Include>
      <Arn>string</Arn>
    </Include>
  </SelectionCriteria>
</StorageLensGroupLevel>
</AccountLevel>
<Include>
  <Buckets>
    <Arn>string</Arn>
  </Buckets>
  <Regions>
    <Region>string</Region>
  </Regions>
</Include>
<Exclude>
  <Buckets>
    <Arn>string</Arn>
  </Buckets>
  <Regions>
    <Region>string</Region>
  </Regions>
</Exclude>
<DataExport>

```

```

<CloudWatchMetrics>
  <IsEnabled>boolean</IsEnabled>
</CloudWatchMetrics>
<S3BucketDestination>
  <AccountId>string</AccountId>
  <Arn>string</Arn>
  <Encryption>
    <SSE-KMS>
      <KeyId>string</KeyId>
    </SSE-KMS>
    <SSE-S3>
    </SSE-S3>
  </Encryption>
  <Format>string</Format>
  <OutputSchemaVersion>string</OutputSchemaVersion>
  <Prefix>string</Prefix>
</S3BucketDestination>
<StorageLensTableDestination>
  <Encryption>
    <SSE-KMS>
      <KeyId>string</KeyId>
    </SSE-KMS>
    <SSE-S3>
    </SSE-S3>
  </Encryption>
  <IsEnabled>boolean</IsEnabled>
</StorageLensTableDestination>
</DataExport>
<ExpandedPrefixesDataExport>
  <S3BucketDestination>
    <AccountId>string</AccountId>
    <Arn>string</Arn>
    <Encryption>
      <SSE-KMS>
        <KeyId>string</KeyId>
      </SSE-KMS>
      <SSE-S3>
      </SSE-S3>
    </Encryption>
    <Format>string</Format>
    <OutputSchemaVersion>string</OutputSchemaVersion>
    <Prefix>string</Prefix>
  </S3BucketDestination>
  <StorageLensTableDestination>

```

```

    <Encryption>
      <SSE-KMS>
        <KeyId>string</KeyId>
      </SSE-KMS>
    <SSE-S3>
    </SSE-S3>
  </Encryption>
  <IsEnabled>boolean</IsEnabled>
</StorageLensTableDestination>
</ExpandedPrefixesDataExport>
<IsEnabled>boolean</IsEnabled>
<AwsOrg>
  <Arn>string</Arn>
</AwsOrg>
<StorageLensArn>string</StorageLensArn>
<PrefixDelimiter>string</PrefixDelimiter>
</StorageLensConfiguration>

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

StorageLensConfiguration

Root level tag for the StorageLensConfiguration parameters.

Required: Yes

AccountLevel

A container for all the account-level configurations of your S3 Storage Lens configuration.

Type: [AccountLevel](#) data type

AwsOrg

A container for the AWS organization for this S3 Storage Lens configuration.

Type: [StorageLensAwsOrg](#) data type

DataExport

A container to specify the properties of your S3 Storage Lens metrics export including, the destination, schema and format.

Type: [StorageLensDataExport](#) data type

[Exclude](#)

A container for what is excluded in this configuration. This container can only be valid if there is no `Include` container submitted, and it's not empty.

Type: [Exclude](#) data type

[ExpandedPrefixesDataExport](#)

A container that configures your S3 Storage Lens expanded prefixes metrics report.

Type: [StorageLensExpandedPrefixesDataExport](#) data type

[Id](#)

A container for the Amazon S3 Storage Lens configuration ID.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\_\.\ ]+`

[Include](#)

A container for what is included in this configuration. This container can only be valid if there is no `Exclude` container submitted, and it's not empty.

Type: [Include](#) data type

[IsEnabled](#)

A container for whether the S3 Storage Lens configuration is enabled.

Type: Boolean

[PrefixDelimiter](#)

A container for all prefix delimiters that are used for object keys in this S3 Storage Lens configuration. The prefix delimiters determine how S3 Storage Lens counts prefix depth, by separating the hierarchical levels in object keys.

Note

- If either a prefix delimiter or existing delimiter is undefined, Amazon S3 uses the delimiter that's defined.

- If both the prefix delimiter and existing delimiter are undefined, S3 uses / as the default delimiter.
- When custom delimiters are used, both the prefix delimiter and existing delimiter must specify the same special character. Otherwise, your request results in an error.

Type: String

Length Constraints: Maximum length of 1.

StorageLensArn

The Amazon Resource Name (ARN) of the S3 Storage Lens configuration. This property is read-only and follows the following format: `arn:aws:s3:us-east-1:example-account-id:storage-lens/your-dashboard-name`

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `arn:[a-z\-\-]+:s3:[a-z0-9\-\-]+:\d{12}:storage\-\lens\/.*`

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetStorageLensConfigurationTagging

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Gets the tags of Amazon S3 Storage Lens configuration. For more information about S3 Storage Lens, see [Assessing your storage activity and usage with Amazon S3 Storage Lens](#) in the *Amazon S3 User Guide*.

Note

To use this action, you must have permission to perform the `s3:GetStorageLensConfigurationTagging` action. For more information, see [Setting permissions to use Amazon S3 Storage Lens](#) in the *Amazon S3 User Guide*.

Request Syntax

```
GET /v20180820/storageLens/storageLensId/tagging HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

storageLensId

The ID of the Amazon S3 Storage Lens configuration.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\_\.\ ]+`

Required: Yes

x-amz-account-id

The account ID of the requester.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<GetStorageLensConfigurationTaggingResult>
  <Tags>
    <Tag>
      <Key>string</Key>
      <Value>string</Value>
    </Tag>
  </Tags>
</GetStorageLensConfigurationTaggingResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[GetStorageLensConfigurationTaggingResult](#)

Root level tag for the `GetStorageLensConfigurationTaggingResult` parameters.

Required: Yes

[Tags](#)

The tags of S3 Storage Lens configuration requested.

Type: Array of [StorageLensTag](#) data types

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetStorageLensGroup

Service: Amazon S3 Control

Retrieves the Storage Lens group configuration details.

To use this operation, you must have the permission to perform the `s3:GetStorageLensGroup` action. For more information about the required Storage Lens Groups permissions, see [Setting account permissions to use S3 Storage Lens groups](#).

For information about Storage Lens groups errors, see [List of Amazon S3 Storage Lens error codes](#).

Request Syntax

```
GET /v20180820/storagegroup/name HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the Storage Lens group that you're trying to retrieve the configuration details for.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\_\+]`

Required: Yes

x-amz-account-id

The AWS account ID associated with the Storage Lens group that you're trying to retrieve the details for.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<StorageLensGroup>
  <Name>string</Name>
  <Filter>
    <And>
      <MatchAnyPrefix>
        <Prefix>string</Prefix>
      </MatchAnyPrefix>
      <MatchAnySuffix>
        <Suffix>string</Suffix>
      </MatchAnySuffix>
      <MatchAnyTag>
        <Tag>
          <Key>string</Key>
          <Value>string</Value>
        </Tag>
      </MatchAnyTag>
      <MatchObjectAge>
        <DaysGreaterThan>integer</DaysGreaterThan>
        <DaysLessThan>integer</DaysLessThan>
      </MatchObjectAge>
      <MatchObjectSize>
        <BytesGreaterThan>long</BytesGreaterThan>
        <BytesLessThan>long</BytesLessThan>
      </MatchObjectSize>
    </And>
  </MatchAnyPrefix>
    <Prefix>string</Prefix>
  </MatchAnyPrefix>
  <MatchAnySuffix>
    <Suffix>string</Suffix>
  </MatchAnySuffix>
  <MatchAnyTag>
    <Tag>
      <Key>string</Key>
      <Value>string</Value>
```

```

    </Tag>
  </MatchAnyTag>
  <MatchObjectAge>
    <DaysGreaterThan>integer</DaysGreaterThan>
    <DaysLessThan>integer</DaysLessThan>
  </MatchObjectAge>
  <MatchObjectSize>
    <BytesGreaterThan>long</BytesGreaterThan>
    <BytesLessThan>long</BytesLessThan>
  </MatchObjectSize>
  <Or>
    <MatchAnyPrefix>
      <Prefix>string</Prefix>
    </MatchAnyPrefix>
    <MatchAnySuffix>
      <Suffix>string</Suffix>
    </MatchAnySuffix>
    <MatchAnyTag>
      <Tag>
        <Key>string</Key>
        <Value>string</Value>
      </Tag>
    </MatchAnyTag>
    <MatchObjectAge>
      <DaysGreaterThan>integer</DaysGreaterThan>
      <DaysLessThan>integer</DaysLessThan>
    </MatchObjectAge>
    <MatchObjectSize>
      <BytesGreaterThan>long</BytesGreaterThan>
      <BytesLessThan>long</BytesLessThan>
    </MatchObjectSize>
  </Or>
</Filter>
<StorageLensGroupArn>string</StorageLensGroupArn>
</StorageLensGroup>

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[StorageLensGroup](#)

Root level tag for the StorageLensGroup parameters.

Required: Yes

[Filter](#)

Sets the criteria for the Storage Lens group data that is displayed. For multiple filter conditions, the AND or OR logical operator is used.

Type: [StorageLensGroupFilter](#) data type

[Name](#)

Contains the name of the Storage Lens group.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [a-zA-Z0-9\-__]+

[StorageLensGroupArn](#)

Contains the Amazon Resource Name (ARN) of the Storage Lens group. This property is read-only.

Type: String

Length Constraints: Minimum length of 4. Maximum length of 1024.

Pattern: arn:[a-z\-\-]+:s3:[a-z0-9\-\-]+:\d{12}:storage\-\lens\-\group\/*

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListAccessGrants

Service: Amazon S3 Control

Returns the list of access grants in your S3 Access Grants instance.

Permissions

You must have the `s3:ListAccessGrants` permission to use this operation.

Request Syntax

```
GET /v20180820/accessgrantsinstance/grants?
application_arn=ApplicationArn&granteeidentifier=GranteeIdentifier&granteetype=GranteeType&gran
HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

application_arn

The Amazon Resource Name (ARN) of an AWS IAM Identity Center application associated with your Identity Center instance. If the grant includes an application ARN, the grantee can only access the S3 data through this application.

Length Constraints: Minimum length of 10. Maximum length of 1224.

Pattern: `arn:[^:]+:sso::\d{12}:application/.*$`

granteeidentifier

The unique identifier of the Grantee. If the grantee type is IAM, the identifier is the IAM Amazon Resource Name (ARN) of the user or role. If the grantee type is a directory user or group, the identifier is 128-bit universally unique identifier (UUID) in the format `a1b2c3d4-5678-90ab-cdef-EXAMPLE11111`. You can obtain this UUID from your AWS IAM Identity Center instance.

granteetype

The type of the grantee to which access has been granted. It can be one of the following values:

- **IAM** - An IAM user or role.
- **DIRECTORY_USER** - Your corporate directory user. You can use this option if you have added your corporate identity directory to IAM Identity Center and associated the IAM Identity Center instance with your S3 Access Grants instance.
- **DIRECTORY_GROUP** - Your corporate directory group. You can use this option if you have added your corporate identity directory to IAM Identity Center and associated the IAM Identity Center instance with your S3 Access Grants instance.

Valid Values: **DIRECTORY_USER** | **DIRECTORY_GROUP** | **IAM**

grantscope

The S3 path of the data to which you are granting access. It is the result of appending the `Subprefix` to the location scope.

Length Constraints: Minimum length of 1. Maximum length of 2000.

Pattern: `^.+`

maxResults

The maximum number of access grants that you would like returned in the `List Access Grants` response. If the results include the pagination token `NextToken`, make another call using the `NextToken` to determine if there are more results.

Valid Range: Minimum value of 0. Maximum value of 1000.

nextToken

A pagination token to request the next page of results. Pass this value into a subsequent `List Access Grants` request in order to retrieve the next page of results.

permission

The type of permission granted to your S3 data, which can be set to one of the following values:

- **READ** – Grant read-only access to the S3 data.
- **WRITE** – Grant write-only access to the S3 data.
- **READWRITE** – Grant both read and write access to the S3 data.

Valid Values: **READ** | **WRITE** | **READWRITE**

x-amz-account-id

The AWS account ID of the S3 Access Grants instance.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<ListAccessGrantsResult>
  <NextToken>string</NextToken>
  <AccessGrantsList>
    <AccessGrant>
      <AccessGrantArn>string</AccessGrantArn>
      <AccessGrantId>string</AccessGrantId>
      <AccessGrantsLocationConfiguration>
        <S3SubPrefix>string</S3SubPrefix>
      </AccessGrantsLocationConfiguration>
      <AccessGrantsLocationId>string</AccessGrantsLocationId>
      <ApplicationArn>string</ApplicationArn>
      <CreatedAt>timestamp</CreatedAt>
      <Grantee>
        <GranteeIdentifier>string</GranteeIdentifier>
        <GranteeType>string</GranteeType>
      </Grantee>
      <GrantScope>string</GrantScope>
      <Permission>string</Permission>
    </AccessGrant>
  </AccessGrantsList>
</ListAccessGrantsResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

ListAccessGrantsResult

Root level tag for the ListAccessGrantsResult parameters.

Required: Yes

AccessGrantsList

A container for a list of grants in an S3 Access Grants instance.

Type: Array of [ListAccessGrantEntry](#) data types

NextToken

A pagination token to request the next page of results. Pass this value into a subsequent List Access Grants request in order to retrieve the next page of results.

Type: String

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListAccessGrantsInstances

Service: Amazon S3 Control

Returns a list of S3 Access Grants instances. An S3 Access Grants instance serves as a logical grouping for your individual access grants. You can only have one S3 Access Grants instance per Region per account.

Permissions

You must have the `s3:ListAccessGrantsInstances` permission to use this operation.

Request Syntax

```
GET /v20180820/accessgrantsinstances?maxResults=MaxResults&nextToken=NextToken HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

maxResults

The maximum number of access grants that you would like returned in the `List Access Grants` response. If the results include the pagination token `NextToken`, make another call using the `NextToken` to determine if there are more results.

Valid Range: Minimum value of 0. Maximum value of 1000.

nextToken

A pagination token to request the next page of results. Pass this value into a subsequent `List Access Grants Instances` request in order to retrieve the next page of results.

x-amz-account-id

The AWS account ID of the S3 Access Grants instance.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<ListAccessGrantsInstancesResult>
  <NextToken>string</NextToken>
  <AccessGrantsInstancesList>
    <AccessGrantsInstance>
      <AccessGrantsInstanceArn>string</AccessGrantsInstanceArn>
      <AccessGrantsInstanceId>string</AccessGrantsInstanceId>
      <CreatedAt>timestamp</CreatedAt>
      <IdentityCenterApplicationArn>string</IdentityCenterApplicationArn>
      <IdentityCenterArn>string</IdentityCenterArn>
      <IdentityCenterInstanceArn>string</IdentityCenterInstanceArn>
    </AccessGrantsInstance>
  </AccessGrantsInstancesList>
</ListAccessGrantsInstancesResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[ListAccessGrantsInstancesResult](#)

Root level tag for the ListAccessGrantsInstancesResult parameters.

Required: Yes

[AccessGrantsInstancesList](#)

A container for a list of S3 Access Grants instances.

Type: Array of [ListAccessGrantsInstanceEntry](#) data types

[NextToken](#)

A pagination token to request the next page of results. Pass this value into a subsequent List Access Grants Instances request in order to retrieve the next page of results.

Type: String

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListAccessGrantsLocations

Service: Amazon S3 Control

Returns a list of the locations registered in your S3 Access Grants instance.

Permissions

You must have the `s3:ListAccessGrantsLocations` permission to use this operation.

Request Syntax

```
GET /v20180820/accessgrantsinstance/locations?
locationscope=LocationScope&maxResults=MaxResults&nextToken=NextToken HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

locationscope

The S3 path to the location that you are registering. The location scope can be the default S3 location `s3://`, the S3 path to a bucket `s3://<bucket>`, or the S3 path to a bucket and prefix `s3://<bucket>/<prefix>`. A prefix in S3 is a string of characters at the beginning of an object key name used to organize the objects that you store in your S3 buckets. For example, object key names that start with the `engineering/` prefix or object key names that start with the `marketing/campaigns/` prefix.

Length Constraints: Minimum length of 1. Maximum length of 2000.

Pattern: `^.+`

maxResults

The maximum number of access grants that you would like returned in the `List Access Grants` response. If the results include the pagination token `NextToken`, make another call using the `NextToken` to determine if there are more results.

Valid Range: Minimum value of 0. Maximum value of 1000.

nextToken

A pagination token to request the next page of results. Pass this value into a subsequent `List Access Grants Locations` request in order to retrieve the next page of results.

x-amz-account-id

The AWS account ID of the S3 Access Grants instance.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<ListAccessGrantsLocationsResult>
  <NextToken>string</NextToken>
  <AccessGrantsLocationsList>
    <AccessGrantsLocation>
      <AccessGrantsLocationArn>string</AccessGrantsLocationArn>
      <AccessGrantsLocationId>string</AccessGrantsLocationId>
      <CreatedAt>timestamp</CreatedAt>
      <IAMRoleArn>string</IAMRoleArn>
      <LocationScope>string</LocationScope>
    </AccessGrantsLocation>
  </AccessGrantsLocationsList>
</ListAccessGrantsLocationsResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[ListAccessGrantsLocationsResult](#)

Root level tag for the ListAccessGrantsLocationsResult parameters.

Required: Yes

[AccessGrantsLocationsList](#)

A container for a list of registered locations in an S3 Access Grants instance.

Type: Array of [ListAccessGrantsLocationsEntry](#) data types

[NextToken](#)

A pagination token to request the next page of results. Pass this value into a subsequent List Access Grants Locations request in order to retrieve the next page of results.

Type: String

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListAccessPoints

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Returns a list of the access points. You can retrieve up to 1,000 access points per call. If the call returns more than 1,000 access points (or the number specified in `maxResults`, whichever is less), the response will include a continuation token that you can use to list the additional access points.

Returns only access points attached to S3 buckets by default. To return all access points specify `DataSourceType` as `ALL`.

All Amazon S3 on Outposts REST API requests for this action require an additional parameter of `x-amz-outpost-id` to be passed with the request. In addition, you must use an S3 on Outposts endpoint hostname prefix instead of `s3-control`. For an example of the request syntax for Amazon S3 on Outposts that uses the S3 on Outposts endpoint hostname prefix and the `x-amz-outpost-id` derived by using the access point ARN, see the [Examples](#) section.

The following actions are related to `ListAccessPoints`:

- [CreateAccessPoint](#)
- [DeleteAccessPoint](#)
- [GetAccessPoint](#)

Request Syntax

```
GET /v20180820/accesspoint?
bucket=Bucket&dataSourceId=DataSourceId&dataSourceType=DataSourceType&maxResults=MaxResults&nextToken=NextToken
HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

bucket

The name of the bucket whose associated access points you want to list.

For using this parameter with Amazon S3 on Outposts with the REST API, you must specify the name and the `x-amz-outpost-id` as well.

For using this parameter with S3 on Outposts with the AWS SDK and CLI, you must specify the ARN of the bucket accessed in the format `arn:aws:s3-outposts:<Region>:<account-id>:outpost/<outpost-id>/bucket/<my-bucket-name>`. For example, to access the bucket `reports` through Outpost `my-outpost` owned by account `123456789012` in Region `us-west-2`, use the URL encoding of `arn:aws:s3-outposts:us-west-2:123456789012:outpost/my-outpost/bucket/reports`. The value must be URL encoded.

Length Constraints: Minimum length of 3. Maximum length of 255.

dataSourceId

The unique identifier for the data source of the access point.

Length Constraints: Maximum length of 191.

dataSourceType

The type of the data source that the access point is attached to. Returns only access points attached to S3 buckets by default. To return all access points specify `DataSourceType` as `ALL`.

maxResults

The maximum number of access points that you want to include in the list. If the specified bucket has more than this number of access points, then the response will include a continuation token in the `NextToken` field that you can use to retrieve the next page of access points.

Valid Range: Minimum value of 0. Maximum value of 1000.

nextToken

A continuation token. If a previous call to `ListAccessPoints` returned a continuation token in the `NextToken` field, then providing that value here causes Amazon S3 to retrieve the next page of results.

Length Constraints: Minimum length of 1. Maximum length of 1024.

x-amz-account-id

The AWS account ID for the account that owns the specified access points.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<ListAccessPointsResult>
  <AccessPointList>
    <AccessPoint>
      <AccessPointArn>string</AccessPointArn>
      <Alias>string</Alias>
      <Bucket>string</Bucket>
      <BucketAccountId>string</BucketAccountId>
      <DataSourceId>string</DataSourceId>
      <DataSourceType>string</DataSourceType>
      <Name>string</Name>
      <NetworkOrigin>string</NetworkOrigin>
      <VpcConfiguration>
        <VpcId>string</VpcId>
      </VpcConfiguration>
    </AccessPoint>
  </AccessPointList>
  <NextToken>string</NextToken>
</ListAccessPointsResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

ListAccessPointsResult

Root level tag for the ListAccessPointsResult parameters.

Required: Yes

AccessPointList

Contains identification and configuration information for one or more access points associated with the specified bucket.

Type: Array of [AccessPoint](#) data types

NextToken

If the specified bucket has more access points than can be returned in one call to this API, this field contains a continuation token that you can provide in subsequent calls to this API to retrieve additional access points.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Examples

Sample request syntax for ListAccessPoints for Amazon S3 on Outposts

The following request returns a list access points of the specified Amazon S3 on Outposts bucket `example-outpost-bucket`.

```
GET /v20180820/accesspoint?Bucket=example-outpost-  
bucket&MaxResults=MaxResults&NextToken=NextToken HTTP/1.1  
Host: s3-outposts.<Region>.amazonaws.com  
Date: Wed, 28 Oct 2020 22:32:00 GMT  
Authorization: authorization string  
x-amz-account-id: example-account-id  
x-amz-outpost-id: op-01ac5d28a6a232904
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListAccessPointsForDirectoryBuckets

Service: Amazon S3 Control

Returns a list of the access points that are owned by the AWS account and that are associated with the specified directory bucket.

To list access points for general purpose buckets, see [ListAccesspoints](#).

To use this operation, you must have the permission to perform the `s3express:ListAccessPointsForDirectoryBuckets` action.

For information about REST API errors, see [REST error responses](#).

Request Syntax

```
GET /v20180820/accesspointfordirectory?
directoryBucket=DirectoryBucket&maxResults=MaxResults&nextToken=NextToken HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

[directoryBucket](#)

The name of the directory bucket associated with the access points you want to list.

Length Constraints: Minimum length of 3. Maximum length of 255.

[maxResults](#)

The maximum number of access points that you would like returned in the `ListAccessPointsForDirectoryBuckets` response. If the directory bucket is associated with more than this number of access points, the results include the pagination token `NextToken`. Make another call using the `NextToken` to retrieve more results.

Valid Range: Minimum value of 0. Maximum value of 1000.

[nextToken](#)

If `NextToken` is returned, there are more access points available than requested in the `maxResults` value. The value of `NextToken` is a unique pagination token for each page. Make

the call again using the returned token to retrieve the next page. Keep all other arguments unchanged. Each pagination token expires after 24 hours.

Length Constraints: Minimum length of 1. Maximum length of 1024.

[x-amz-account-id](#)

The AWS account ID that owns the access points.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<ListAccessPointsForDirectoryBucketsResult>
  <AccessPointList>
    <AccessPoint>
      <AccessPointArn>string</AccessPointArn>
      <Alias>string</Alias>
      <Bucket>string</Bucket>
      <BucketAccountId>string</BucketAccountId>
      <DataSourceId>string</DataSourceId>
      <DataSourceType>string</DataSourceType>
      <Name>string</Name>
      <NetworkOrigin>string</NetworkOrigin>
      <VpcConfiguration>
        <VpcId>string</VpcId>
      </VpcConfiguration>
    </AccessPoint>
  </AccessPointList>
  <NextToken>string</NextToken>
</ListAccessPointsForDirectoryBucketsResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[ListAccessPointsForDirectoryBucketsResult](#)

Root level tag for the ListAccessPointsForDirectoryBucketsResult parameters.

Required: Yes

[AccessPointList](#)

Contains identification and configuration information for one or more access points associated with the directory bucket.

Type: Array of [AccessPoint](#) data types

[NextToken](#)

If NextToken is returned, there are more access points available than requested in the maxResults value. The value of NextToken is a unique pagination token for each page. Make the call again using the returned token to retrieve the next page. Keep all other arguments unchanged. Each pagination token expires after 24 hours.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)

- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListAccessPointsForObjectLambda

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Returns some or all (up to 1,000) access points associated with the Object Lambda Access Point per call. If there are more access points than what can be returned in one call, the response will include a continuation token that you can use to list the additional access points.

The following actions are related to ListAccessPointsForObjectLambda:

- [CreateAccessPointForObjectLambda](#)
- [DeleteAccessPointForObjectLambda](#)
- [GetAccessPointForObjectLambda](#)

Request Syntax

```
GET /v20180820/accesspointforobjectlambda?maxResults=MaxResults&nextToken=NextToken
HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

maxResults

The maximum number of access points that you want to include in the list. The response may contain fewer access points but will never contain more. If there are more than this number of access points, then the response will include a continuation token in the NextToken field that you can use to retrieve the next page of access points.

Valid Range: Minimum value of 0. Maximum value of 1000.

nextToken

If the list has more access points than can be returned in one call to this API, this field contains a continuation token that you can provide in subsequent calls to this API to retrieve additional access points.

Length Constraints: Minimum length of 1. Maximum length of 1024.

x-amz-account-id

The account ID for the account that owns the specified Object Lambda Access Point.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<ListAccessPointsForObjectLambdaResult>
  <ObjectLambdaAccessPointList>
    <ObjectLambdaAccessPoint>
      <Alias>
        <Status>string</Status>
        <Value>string</Value>
      </Alias>
      <Name>string</Name>
      <ObjectLambdaAccessPointArn>string</ObjectLambdaAccessPointArn>
    </ObjectLambdaAccessPoint>
  </ObjectLambdaAccessPointList>
  <NextToken>string</NextToken>
</ListAccessPointsForObjectLambdaResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[ListAccessPointsForObjectLambdaResult](#)

Root level tag for the ListAccessPointsForObjectLambdaResult parameters.

Required: Yes

[NextToken](#)

If the list has more access points than can be returned in one call to this API, this field contains a continuation token that you can provide in subsequent calls to this API to retrieve additional access points.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

[ObjectLambdaAccessPointList](#)

Returns list of Object Lambda Access Points.

Type: Array of [ObjectLambdaAccessPoint](#) data types

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListCallerAccessGrants

Service: Amazon S3 Control

Use this API to list the access grants that grant the caller access to Amazon S3 data through S3 Access Grants. The caller (grantee) can be an AWS Identity and Access Management (IAM) identity or AWS Identity Center corporate directory identity. You must pass the AWS account of the S3 data owner (grantor) in the request. You can, optionally, narrow the results by `GrantScope`, using a fragment of the data's S3 path, and S3 Access Grants will return only the grants with a path that contains the path fragment. You can also pass the `AllowedByApplication` filter in the request, which returns only the grants authorized for applications, whether the application is the caller's Identity Center application or any other application (ALL). For more information, see [List the caller's access grants](#) in the *Amazon S3 User Guide*.

Permissions

You must have the `s3:ListCallerAccessGrants` permission to use this operation.

Request Syntax

```
GET /v20180820/accessgrantsinstance/caller/grants?
allowedByApplication=AllowedByApplication&grantscope=GrantScope&maxResults=MaxResults&nextToken=
HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

[allowedByApplication](#)

If this optional parameter is passed in the request, a filter is applied to the results. The results will include only the access grants for the caller's Identity Center application or for any other applications (ALL).

[grantscope](#)

The S3 path of the data that you would like to access. Must start with `s3://`. You can optionally pass only the beginning characters of a path, and S3 Access Grants will search for all applicable grants for the path fragment.

Length Constraints: Minimum length of 1. Maximum length of 2000.

Pattern: `^\.+`

maxResults

The maximum number of access grants that you would like returned in the `List Caller Access Grants` response. If the results include the pagination token `NextToken`, make another call using the `NextToken` to determine if there are more results.

Valid Range: Minimum value of 0. Maximum value of 1000.

nextToken

A pagination token to request the next page of results. Pass this value into a subsequent `List Caller Access Grants` request in order to retrieve the next page of results.

x-amz-account-id

The AWS account ID of the S3 Access Grants instance.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<ListCallerAccessGrantsResult>
  <NextToken>string</NextToken>
  <CallerAccessGrantsList>
    <AccessGrant>
      <ApplicationArn>string</ApplicationArn>
      <GrantScope>string</GrantScope>
      <Permission>string</Permission>
    </AccessGrant>
  </CallerAccessGrantsList>
```

```
</ListCallerAccessGrantsResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[ListCallerAccessGrantsResult](#)

Root level tag for the ListCallerAccessGrantsResult parameters.

Required: Yes

[CallerAccessGrantsList](#)

A list of the caller's access grants that were created using S3 Access Grants and that grant the caller access to the S3 data of the AWS account ID that was specified in the request.

Type: Array of [ListCallerAccessGrantsEntry](#) data types

[NextToken](#)

A pagination token that you can use to request the next page of results. Pass this value into a subsequent List Caller Access Grants request in order to retrieve the next page of results.

Type: String

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)

- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListJobs

Service: Amazon S3 Control

Lists current S3 Batch Operations jobs as well as the jobs that have ended within the last 90 days for the AWS account making the request. For more information, see [S3 Batch Operations](#) in the *Amazon S3 User Guide*.

Permissions

To use the ListJobs operation, you must have permission to perform the `s3:ListJobs` action.

Related actions include:

- [CreateJob](#)
- [DescribeJob](#)
- [UpdateJobPriority](#)
- [UpdateJobStatus](#)

Request Syntax

```
GET /v20180820/jobs?jobStatuses=JobStatuses&maxResults=MaxResults&nextToken=NextToken
HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

[jobStatuses](#)

The List Jobs request returns jobs that match the statuses listed in this element.

Valid Values: Active | Cancelled | Cancelling | Complete | Completing
| Failed | Failing | New | Paused | Pausing | Preparing | Ready |
Suspended

maxResults

The maximum number of jobs that Amazon S3 will include in the `ListJobs` response. If there are more jobs than this number, the response will include a pagination token in the `NextToken` field to enable you to retrieve the next page of results.

Valid Range: Minimum value of 0. Maximum value of 1000.

nextToken

A pagination token to request the next page of results. Use the token that Amazon S3 returned in the `NextToken` element of the `ListJobsResult` from the previous `ListJobs` request.

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `^[A-Za-z0-9\+\:\\/\=\?\#\_\-]+$`

x-amz-account-id

The AWS account ID associated with the S3 Batch Operations job.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<ListJobsResult>
  <NextToken>string</NextToken>
  <Jobs>
    <JobListDescriptor>
      <CreationTime>timestamp</CreationTime>
      <Description>string</Description>
      <JobId>string</JobId>
```

```

<Operation>string</Operation>
<Priority>integer</Priority>
<ProgressSummary>
  <NumberOfTasksFailed>long</NumberOfTasksFailed>
  <NumberOfTasksSucceeded>long</NumberOfTasksSucceeded>
  <Timers>
    <ElapsedTimeInActiveSeconds>long</ElapsedTimeInActiveSeconds>
  </Timers>
  <TotalNumberOfTasks>long</TotalNumberOfTasks>
</ProgressSummary>
<Status>string</Status>
<TerminationDate>timestamp</TerminationDate>
</JobListDescriptor>
</Jobs>
</ListJobsResult>

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[ListJobsResult](#)

Root level tag for the ListJobsResult parameters.

Required: Yes

[Jobs](#)

The list of current jobs and jobs that have ended within the last 30 days.

Type: Array of [JobListDescriptor](#) data types

[NextToken](#)

If the List Jobs request produced more than the maximum number of results, you can pass this value into a subsequent List Jobs request in order to retrieve the next page of results.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `^[A-Za-z0-9\+\:\.\=\?\#\_\-]+$`

Errors

InternalServiceException

HTTP Status Code: 500

InvalidNextTokenException

HTTP Status Code: 400

InvalidRequestException

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListMultiRegionAccessPoints

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Returns a list of the Multi-Region Access Points currently associated with the specified AWS account. Each call can return up to 100 Multi-Region Access Points, the maximum number of Multi-Region Access Points that can be associated with a single account.

This action will always be routed to the US West (Oregon) Region. For more information about the restrictions around working with Multi-Region Access Points, see [Multi-Region Access Point restrictions and limitations](#) in the *Amazon S3 User Guide*.

The following actions are related to ListMultiRegionAccessPoint:

- [CreateMultiRegionAccessPoint](#)
- [DeleteMultiRegionAccessPoint](#)
- [DescribeMultiRegionAccessPointOperation](#)
- [GetMultiRegionAccessPoint](#)

Request Syntax

```
GET /v20180820/mrap/instances?maxResults=MaxResults&nextToken=NextToken HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

[maxResults](#)

Not currently used. Do not use this parameter.

Valid Range: Minimum value of 0. Maximum value of 1000.

nextToken

Not currently used. Do not use this parameter.

Length Constraints: Minimum length of 1. Maximum length of 1024.

x-amz-account-id

The AWS account ID for the owner of the Multi-Region Access Point.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<ListMultiRegionAccessPointsResult>
  <AccessPoints>
    <AccessPoint>
      <Alias>string</Alias>
      <CreatedAt>timestamp</CreatedAt>
      <Name>string</Name>
      <PublicAccessBlock>
        <BlockPublicAcls>boolean</BlockPublicAcls>
        <BlockPublicPolicy>boolean</BlockPublicPolicy>
        <IgnorePublicAcls>boolean</IgnorePublicAcls>
        <RestrictPublicBuckets>boolean</RestrictPublicBuckets>
      </PublicAccessBlock>
    <Regions>
      <Region>
        <Bucket>string</Bucket>
        <BucketAccountId>string</BucketAccountId>
        <Region>string</Region>
      </Region>
    </Regions>
  </AccessPoints>
</ListMultiRegionAccessPointsResult>
```

```
<Status>string</Status>
</AccessPoint>
</AccessPoints>
<NextToken>string</NextToken>
</ListMultiRegionAccessPointsResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[ListMultiRegionAccessPointsResult](#)

Root level tag for the ListMultiRegionAccessPointsResult parameters.

Required: Yes

[AccessPoints](#)

The list of Multi-Region Access Points associated with the user.

Type: Array of [MultiRegionAccessPointReport](#) data types

[NextToken](#)

If the specified bucket has more Multi-Region Access Points than can be returned in one call to this action, this field contains a continuation token. You can use this token in subsequent calls to this action to retrieve additional Multi-Region Access Points.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)

- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListRegionalBuckets

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Returns a list of all Outposts buckets in an Outpost that are owned by the authenticated sender of the request. For more information, see [Using Amazon S3 on Outposts](#) in the *Amazon S3 User Guide*.

For an example of the request syntax for Amazon S3 on Outposts that uses the S3 on Outposts endpoint hostname prefix and `x-amz-outpost-id` in your request, see the [Examples](#) section.

Request Syntax

```
GET /v20180820/bucket?maxResults=MaxResults&nextToken=NextToken HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
x-amz-outpost-id: OutpostId
```

URI Request Parameters

The request uses the following URI parameters.

[maxResults](#)

Valid Range: Minimum value of 0. Maximum value of 1000.

[nextToken](#)

Length Constraints: Minimum length of 1. Maximum length of 1024.

[x-amz-account-id](#)

The AWS account ID of the Outposts bucket.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

x-amz-outpost-id

The ID of the AWS Outposts resource.

Note

This ID is required by Amazon S3 on Outposts buckets.

Length Constraints: Minimum length of 1. Maximum length of 64.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<ListRegionalBucketsResult>
  <RegionalBucketList>
    <RegionalBucket>
      <Bucket>string</Bucket>
      <BucketArn>string</BucketArn>
      <CreationDate>timestamp</CreationDate>
      <OutpostId>string</OutpostId>
      <PublicAccessBlockEnabled>boolean</PublicAccessBlockEnabled>
    </RegionalBucket>
  </RegionalBucketList>
  <NextToken>string</NextToken>
</ListRegionalBucketsResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

ListRegionalBucketsResult

Root level tag for the ListRegionalBucketsResult parameters.

Required: Yes

[NextToken](#)

NextToken is sent when `isTruncated` is true, which means there are more buckets that can be listed. The next list requests to Amazon S3 can be continued with this NextToken. NextToken is obfuscated and is not a real key.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

[RegionalBucketList](#)

Type: Array of [RegionalBucket](#) data types

Examples

Sample request to list an account's Outposts buckets

This request lists regional buckets.

```
GET /v20180820/bucket HTTP /1.1
Host:s3-outposts.us-west-2.amazonaws.com
Content-Length: 0
x-amz-outpost-id: op-01ac5d28a6a232904
x-amz-account-id: example-account-id
Date: Wed, 01 Mar 2006 12:00:00 GMT
Authorization: authorization string
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListStorageLensConfigurations

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Gets a list of Amazon S3 Storage Lens configurations. For more information about S3 Storage Lens, see [Assessing your storage activity and usage with Amazon S3 Storage Lens](#) in the *Amazon S3 User Guide*.

Note

To use this action, you must have permission to perform the `s3:ListStorageLensConfigurations` action. For more information, see [Setting permissions to use Amazon S3 Storage Lens](#) in the *Amazon S3 User Guide*.

Request Syntax

```
GET /v20180820/storagelens?nextToken=NextToken HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

nextToken

A pagination token to request the next page of results.

x-amz-account-id

The account ID of the requester.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<ListStorageLensConfigurationsResult>
  <NextToken>string</NextToken>
  <StorageLensConfiguration>
    <HomeRegion>string</HomeRegion>
    <Id>string</Id>
    <IsEnabled>boolean</IsEnabled>
    <StorageLensArn>string</StorageLensArn>
  </StorageLensConfiguration>
  ...
</ListStorageLensConfigurationsResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

ListStorageLensConfigurationsResult

Root level tag for the ListStorageLensConfigurationsResult parameters.

Required: Yes

NextToken

If the request produced more than the maximum number of S3 Storage Lens configuration results, you can pass this value into a subsequent request to retrieve the next page of results.

Type: String

StorageLensConfiguration

A list of S3 Storage Lens configurations.

Type: Array of [ListStorageLensConfigurationEntry](#) data types

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListStorageLensGroups

Service: Amazon S3 Control

Lists all the Storage Lens groups in the specified home Region.

To use this operation, you must have the permission to perform the `s3:ListStorageLensGroups` action. For more information about the required Storage Lens Groups permissions, see [Setting account permissions to use S3 Storage Lens groups](#).

For information about Storage Lens groups errors, see [List of Amazon S3 Storage Lens error codes](#).

Request Syntax

```
GET /v20180820/storagelensgroup?nextToken=NextToken HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

nextToken

The token for the next set of results, or `null` if there are no more results.

x-amz-account-id

The AWS account ID that owns the Storage Lens groups.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

```
<?xml version="1.0" encoding="UTF-8"?>
<ListStorageLensGroupsResult>
  <NextToken>string</NextToken>
  <StorageLensGroup>
    <HomeRegion>string</HomeRegion>
    <Name>string</Name>
    <StorageLensGroupArn>string</StorageLensGroupArn>
  </StorageLensGroup>
  ...
</ListStorageLensGroupsResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[ListStorageLensGroupsResult](#)

Root level tag for the ListStorageLensGroupsResult parameters.

Required: Yes

[NextToken](#)

If NextToken is returned, there are more Storage Lens groups results available. The value of NextToken is a unique pagination token for each page. Make the call again using the returned token to retrieve the next page. Keep all other arguments unchanged. Each pagination token expires after 24 hours.

Type: String

[StorageLensGroup](#)

The list of Storage Lens groups that exist in the specified home Region.

Type: Array of [ListStorageLensGroupEntry](#) data types

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListTagsForResource

Service: Amazon S3 Control

This operation allows you to list all of the tags for a specified resource. Each tag is a label consisting of a key and value. Tags can help you organize, track costs for, and control access to resources.

Note

This operation is only supported for the following Amazon S3 resources:

- [General purpose buckets](#)
- [Access Points for directory buckets](#)
- [Access Points for general purpose buckets](#)
- [Directory buckets](#)
- [S3 Storage Lens groups](#)
- [S3 Access Grants instances, registered locations, and grants.](#)

Permissions

For general purpose buckets, access points for general purpose buckets, Storage Lens groups, and S3 Access Grants, you must have the `s3:ListTagsForResource` permission to use this operation.

Directory bucket permissions

For directory buckets, you must have the `s3express:ListTagsForResource` permission to use this operation. For more information about directory buckets policies and permissions, see [Identity and Access Management \(IAM\) for S3 Express One Zone](#) in the *Amazon S3 User Guide*.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `s3express-control.region.amazonaws.com`.

For information about S3 Tagging errors, see [List of Amazon S3 Tagging error codes](#).

Request Syntax

```
GET /v20180820/tags/resourceArn+ HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

resourceArn

The Amazon Resource Name (ARN) of the S3 resource that you want to list tags for. The tagged resource can be a directory bucket, S3 Storage Lens group or S3 Access Grants instance, registered location, or grant.

Length Constraints: Maximum length of 1011.

Pattern: `arn:[^:]+:s3(express)?:[^:]*`

Required: Yes

x-amz-account-id

The AWS account ID of the resource owner.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<ListTagsForResourceResult>
  <Tags>
```



```
<Tag>
  <Key>string</Key>
  <Value>string</Value>
</Tag>
</Tags>
</ListTagsForResourceResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

ListTagsForResourceResult

Root level tag for the ListTagsForResourceResult parameters.

Required: Yes

Tags

The AWS resource tags that are associated with the resource.

Type: Array of [Tag](#) data types

Array Members: Minimum number of 0 items. Maximum number of 50 items.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)

- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutAccessGrantsInstanceResourcePolicy

Service: Amazon S3 Control

Updates the resource policy of the S3 Access Grants instance.

Permissions

You must have the `s3:PutAccessGrantsInstanceResourcePolicy` permission to use this operation.

Request Syntax

```
PUT /v20180820/accessgrantsinstance/resourcepolicy HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
<?xml version="1.0" encoding="UTF-8"?>
<PutAccessGrantsInstanceResourcePolicyRequest xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <Policy>string</Policy>
  <Organization>string</Organization>
</PutAccessGrantsInstanceResourcePolicyRequest>
```

URI Request Parameters

The request uses the following URI parameters.

[x-amz-account-id](#)

The AWS account ID of the S3 Access Grants instance.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request accepts the following data in XML format.

PutAccessGrantsInstanceResourcePolicyRequest

Root level tag for the PutAccessGrantsInstanceResourcePolicyRequest parameters.

Required: Yes

Organization

The Organization of the resource policy of the S3 Access Grants instance.

Type: String

Length Constraints: Minimum length of 12. Maximum length of 34.

Pattern: `^[a-z0-9]{10,32}$`

Required: No

Policy

The resource policy of the S3 Access Grants instance that you are updating.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 350000.

Required: Yes

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<PutAccessGrantsInstanceResourcePolicyResult>
  <Policy>string</Policy>
  <Organization>string</Organization>
  <CreatedAt>timestamp</CreatedAt>
</PutAccessGrantsInstanceResourcePolicyResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

PutAccessGrantsInstanceResourcePolicyResult

Root level tag for the PutAccessGrantsInstanceResourcePolicyResult parameters.

Required: Yes

CreatedAt

The date and time when you created the S3 Access Grants instance resource policy.

Type: Timestamp

Organization

The Organization of the resource policy of the S3 Access Grants instance.

Type: String

Length Constraints: Minimum length of 12. Maximum length of 34.

Pattern: `^[a-z0-9]{10,32}$`

Policy

The updated resource policy of the S3 Access Grants instance.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 350000.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)

- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutAccessPointConfigurationForObjectLambda

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Replaces configuration for an Object Lambda Access Point.

The following actions are related to PutAccessPointConfigurationForObjectLambda:

- [GetAccessPointConfigurationForObjectLambda](#)

Request Syntax

```
PUT /v20180820/accesspointforobjectlambda/name/configuration HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
<?xml version="1.0" encoding="UTF-8"?>
<PutAccessPointConfigurationForObjectLambdaRequest xmlns="http://
awss3control.amazonaws.com/doc/2018-08-20/">
  <Configuration>
    <AllowedFeatures>
      <AllowedFeature>string</AllowedFeature>
    </AllowedFeatures>
    <CloudWatchMetricsEnabled>boolean</CloudWatchMetricsEnabled>
    <SupportingAccessPoint>string</SupportingAccessPoint>
    <TransformationConfigurations>
      <TransformationConfiguration>
        <Actions>
          <Action>string</Action>
        </Actions>
        <ContentTransformation>
          <AwsLambda>
            <FunctionArn>string</FunctionArn>
            <FunctionPayload>string</FunctionPayload>
          </AwsLambda>
        </ContentTransformation>
      </TransformationConfiguration>
    </TransformationConfigurations>
  </Configuration>
</PutAccessPointConfigurationForObjectLambdaRequest>
```

```
</Configuration>  
</PutAccessPointConfigurationForObjectLambdaRequest>
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the Object Lambda Access Point.

Length Constraints: Minimum length of 3. Maximum length of 45.

Pattern: `^[a-z0-9]([a-z0-9\-*][a-z0-9])?$`

Required: Yes

x-amz-account-id

The account ID for the account that owns the specified Object Lambda Access Point.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request accepts the following data in XML format.

PutAccessPointConfigurationForObjectLambdaRequest

Root level tag for the PutAccessPointConfigurationForObjectLambdaRequest parameters.

Required: Yes

Configuration

Object Lambda Access Point configuration document.

Type: [ObjectLambdaConfiguration](#) data type

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutAccessPointPolicy

Service: Amazon S3 Control

Associates an access policy with the specified access point. Each access point can have only one policy, so a request made to this API replaces any existing policy associated with the specified access point.

All Amazon S3 on Outposts REST API requests for this action require an additional parameter of `x-amz-outpost-id` to be passed with the request. In addition, you must use an S3 on Outposts endpoint hostname prefix instead of `s3-control`. For an example of the request syntax for Amazon S3 on Outposts that uses the S3 on Outposts endpoint hostname prefix and the `x-amz-outpost-id` derived by using the access point ARN, see the [Examples](#) section.

The following actions are related to `PutAccessPointPolicy`:

- [GetAccessPointPolicy](#)
- [DeleteAccessPointPolicy](#)

Request Syntax

```
PUT /v20180820/accesspoint/name/policy HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
<?xml version="1.0" encoding="UTF-8"?>
<PutAccessPointPolicyRequest xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <Policy>string</Policy>
</PutAccessPointPolicyRequest>
```

URI Request Parameters

The request uses the following URI parameters.

[name](#)

The name of the access point that you want to associate with the specified policy.

For using this parameter with Amazon S3 on Outposts with the REST API, you must specify the name and the `x-amz-outpost-id` as well.

For using this parameter with S3 on Outposts with the AWS SDK and CLI, you must specify the ARN of the access point accessed in the format `arn:aws:s3-outposts:<Region>:<account-id>:outpost/<outpost-id>/accesspoint/<my-accesspoint-name>`. For example, to access the access point `reports-ap` through Outpost `my-outpost` owned by account `123456789012` in Region `us-west-2`, use the URL encoding of `arn:aws:s3-outposts:us-west-2:123456789012:outpost/my-outpost/accesspoint/reports-ap`. The value must be URL encoded.

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

[x-amz-account-id](#)

The AWS account ID for owner of the bucket associated with the specified access point.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request accepts the following data in XML format.

[PutAccessPointPolicyRequest](#)

Root level tag for the `PutAccessPointPolicyRequest` parameters.

Required: Yes

[Policy](#)

The policy that you want to apply to the specified access point. For more information about access point policies, see [Managing data access with Amazon S3 access points](#) or [Managing access to shared datasets in directory buckets with access points](#) in the *Amazon S3 User Guide*.

Type: String

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Sample request syntax for the PutAccessPointPolicy action for Amazon S3 on Outposts access point

This example illustrates one usage of PutAccessPointPolicy.

```
PUT /v20180820/accesspoint/example-access-point/policy HTTP/1.1
Host: s3-outposts.<Region>.amazonaws.com
Date: Wed, 28 Oct 2020 22:32:00 GMT
Authorization: authorization string
x-amz-account-id: example-account-id
x-amz-outpost-id: op-01ac5d28a6a232904
<?xml version="1.0" encoding="UTF-8"?>
    <PutAccessPointPolicyRequest xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
        <Policy>
{
    "Version":"2012-10-17",
    "Id":"AccessPointPolicy-for-example-access-point",
    "Statement":[
        {
            "Sid":"st1",
            "Effect":"Allow",
            "Principal":{"
                "AWS":"example-account-id"
            },
            "Action":"s3-outposts:*",
            "Resource":"arn:aws:s3-outposts:your-Region:example-account-id:outpost/
op-01ac5d28a6a232904/accesspoint/example-access-point
        }
    ]
}
```

```
</Policy>  
</PutAccessPointPolicyRequest>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutAccessPointPolicyForObjectLambda

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Creates or replaces resource policy for an Object Lambda Access Point. For an example policy, see [Creating Object Lambda Access Points](#) in the *Amazon S3 User Guide*.

The following actions are related to PutAccessPointPolicyForObjectLambda:

- [DeleteAccessPointPolicyForObjectLambda](#)
- [GetAccessPointPolicyForObjectLambda](#)

Request Syntax

```
PUT /v20180820/accesspointforobjectlambda/name/policy HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
<?xml version="1.0" encoding="UTF-8"?>
<PutAccessPointPolicyForObjectLambdaRequest xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <Policy>string</Policy>
</PutAccessPointPolicyForObjectLambdaRequest>
```

URI Request Parameters

The request uses the following URI parameters.

[name](#)

The name of the Object Lambda Access Point.

Length Constraints: Minimum length of 3. Maximum length of 45.

Pattern: `^[a-z0-9]([a-z0-9\-*][a-z0-9])?$`

Required: Yes

x-amz-account-id

The account ID for the account that owns the specified Object Lambda Access Point.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request accepts the following data in XML format.

PutAccessPointPolicyForObjectLambdaRequest

Root level tag for the PutAccessPointPolicyForObjectLambdaRequest parameters.

Required: Yes

Policy

Object Lambda Access Point resource policy document.

Type: String

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Sample resource policy

The following illustrates a sample resource policy.

```
{
  "Version" : "2008-10-17",
  "Statement": [{
    "Sid": "Grant account 123456789012 GetObject access",
    "Effect": "Allow",
    "Principal" : {
      "AWS": "arn:aws:iam::123456789012:root"
    },
    "Action": ["s3-object-lambda:GetObject"],
    "Resource": ["arn:aws:s3-object-lambda:us-east-1:123456789012:accesspoint/my-object-lambda-ap"]
  },
  {
    "Sid": "Grant account 444455556666 GetObject access",
    "Effect": "Allow",
    "Principal" : {
      "AWS": "arn:aws:iam::444455556666:root"
    },
    "Action": ["s3-object-lambda:GetObject"],
    "Resource": ["arn:aws:s3-object-lambda:us-east-1:123456789012:accesspoint/my-object-lambda-ap"]
  }
]
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)

- [AWS SDK for Ruby V3](#)

PutAccessPointScope

Service: Amazon S3 Control

Creates or replaces the access point scope for a directory bucket. You can use the access point scope to restrict access to specific prefixes, API operations, or a combination of both.

Note

You can specify any amount of prefixes, but the total length of characters of all prefixes must be less than 256 bytes in size.

To use this operation, you must have the permission to perform the `s3express:PutAccessPointScope` action.

For information about REST API errors, see [REST error responses](#).

Request Syntax

```
PUT /v20180820/accesspoint/name/scope HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
<?xml version="1.0" encoding="UTF-8"?>
<PutAccessPointScopeRequest xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <Scope>
    <Permissions>
      <Permission>string</Permission>
    </Permissions>
    <Prefixes>
      <Prefix>string</Prefix>
    </Prefixes>
  </Scope>
</PutAccessPointScopeRequest>
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the access point with the scope that you want to create or replace.

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

[x-amz-account-id](#)

The AWS account ID that owns the access point with scope that you want to create or replace.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request accepts the following data in XML format.

[PutAccessPointScopeRequest](#)

Root level tag for the PutAccessPointScopeRequest parameters.

Required: Yes

[Scope](#)

Object prefixes, API operations, or a combination of both.

Type: [Scope](#) data type

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutBucketLifecycleConfiguration

Service: Amazon S3 Control

Note

This action puts a lifecycle configuration to an Amazon S3 on Outposts bucket. To put a lifecycle configuration to an S3 bucket, see [PutBucketLifecycleConfiguration](#) in the *Amazon S3 API Reference*.

Creates a new lifecycle configuration for the S3 on Outposts bucket or replaces an existing lifecycle configuration. Outposts buckets only support lifecycle configurations that delete/expire objects after a certain period of time and abort incomplete multipart uploads.

All Amazon S3 on Outposts REST API requests for this action require an additional parameter of `x-amz-outpost-id` to be passed with the request. In addition, you must use an S3 on Outposts endpoint hostname prefix instead of `s3-control`. For an example of the request syntax for Amazon S3 on Outposts that uses the S3 on Outposts endpoint hostname prefix and the `x-amz-outpost-id` derived by using the access point ARN, see the [Examples](#) section.

The following actions are related to `PutBucketLifecycleConfiguration`:

- [GetBucketLifecycleConfiguration](#)
- [DeleteBucketLifecycleConfiguration](#)

Request Syntax

```
PUT /v20180820/bucket/name/lifecycleconfiguration HTTP/1.1
Host: Bucket.s3-control.amazonaws.com
x-amz-account-id: AccountId
<?xml version="1.0" encoding="UTF-8"?>
<LifecycleConfiguration xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <Rules>
    <Rule>
      <AbortIncompleteMultipartUpload>
        <DaysAfterInitiation>integer</DaysAfterInitiation>
      </AbortIncompleteMultipartUpload>
      <Expiration>
        <Date>timestamp</Date>
```

```

    <Days>integer</Days>
    <ExpiredObjectDeleteMarker>boolean</ExpiredObjectDeleteMarker>
  </Expiration>
  <Filter>
    <And>
      <ObjectSizeGreaterThan>long</ObjectSizeGreaterThan>
      <ObjectSizeLessThan>long</ObjectSizeLessThan>
      <Prefix>string</Prefix>
      <Tags>
        <S3Tag>
          <Key>string</Key>
          <Value>string</Value>
        </S3Tag>
      </Tags>
    </And>
    <ObjectSizeGreaterThan>long</ObjectSizeGreaterThan>
    <ObjectSizeLessThan>long</ObjectSizeLessThan>
    <Prefix>string</Prefix>
    <Tag>
      <Key>string</Key>
      <Value>string</Value>
    </Tag>
  </Filter>
  <ID>string</ID>
  <NoncurrentVersionExpiration>
    <NewerNoncurrentVersions>integer</NewerNoncurrentVersions>
    <NoncurrentDays>integer</NoncurrentDays>
  </NoncurrentVersionExpiration>
  <NoncurrentVersionTransitions>
    <NoncurrentVersionTransition>
      <NoncurrentDays>integer</NoncurrentDays>
      <StorageClass>string</StorageClass>
    </NoncurrentVersionTransition>
  </NoncurrentVersionTransitions>
  <Status>string</Status>
  <Transitions>
    <Transition>
      <Date>timestamp</Date>
      <Days>integer</Days>
      <StorageClass>string</StorageClass>
    </Transition>
  </Transitions>
</Rule>
</Rules>

```

```
</LifecycleConfiguration>
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the bucket for which to set the configuration.

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

x-amz-account-id

The AWS account ID of the Outposts bucket.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request accepts the following data in XML format.

LifecycleConfiguration

Root level tag for the LifecycleConfiguration parameters.

Required: Yes

Rules

A lifecycle rule for individual objects in an Outposts bucket.

Type: Array of [LifecycleRule](#) data types

Required: No

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Sample PutBucketLifecycleConfiguration request on an Amazon S3 on Outposts bucket

This request puts a lifecycle configuration on an Outposts bucket named example-outpost-bucket.

```
PUT /v20180820/bucket/example-outpost-bucket/lifecycleconfiguration
HTTP/1.1
Host:s3-outposts.<Region>.amazonaws.com
x-amz-account-id: example-account-id
x-amz-outpost-id: op-01ac5d28a6a232904
Content-Length: 0
Date: Wed, 01 Mar 2006 12:00:00 GMT
Content-MD5: q6yJDlIkBaGGfb3QLY69A==
Authorization: authorization string
Content-Length: 214

<LifecycleConfiguration>
  <Rule>
    <ID>id2</ID>
    <Filter>
      <Prefix>logs/</Prefix>
    </Filter>
    <Status>Enabled</Status>
    <Expiration>
      <Days>365</Days>
    </Expiration>
  </Rule>
</LifecycleConfiguration>
```


See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutBucketPolicy

Service: Amazon S3 Control

Note

This action puts a bucket policy to an Amazon S3 on Outposts bucket. To put a policy on an S3 bucket, see [PutBucketPolicy](#) in the *Amazon S3 API Reference*.

Applies an Amazon S3 bucket policy to an Outposts bucket. For more information, see [Using Amazon S3 on Outposts](#) in the *Amazon S3 User Guide*.

If you are using an identity other than the root user of the AWS account that owns the Outposts bucket, the calling identity must have the PutBucketPolicy permissions on the specified Outposts bucket and belong to the bucket owner's account in order to use this action.

If you don't have PutBucketPolicy permissions, Amazon S3 returns a 403 Access Denied error. If you have the correct permissions, but you're not using an identity that belongs to the bucket owner's account, Amazon S3 returns a 405 Method Not Allowed error.

Important

As a security precaution, the root user of the AWS account that owns a bucket can always use this action, even if the policy explicitly denies the root user the ability to perform this action.

For more information about bucket policies, see [Using Bucket Policies and User Policies](#).

All Amazon S3 on Outposts REST API requests for this action require an additional parameter of `x-amz-outpost-id` to be passed with the request. In addition, you must use an S3 on Outposts endpoint hostname prefix instead of `s3-control`. For an example of the request syntax for Amazon S3 on Outposts that uses the S3 on Outposts endpoint hostname prefix and the `x-amz-outpost-id` derived by using the access point ARN, see the [Examples](#) section.

The following actions are related to PutBucketPolicy:

- [GetBucketPolicy](#)
- [DeleteBucketPolicy](#)

Request Syntax

```
PUT /v20180820/bucket/name/policy HTTP/1.1
Host: Bucket.s3-control.amazonaws.com
x-amz-account-id: AccountId
x-amz-confirm-remove-self-bucket-access: ConfirmRemoveSelfBucketAccess
<?xml version="1.0" encoding="UTF-8"?>
<PutBucketPolicyRequest xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <Policy>string</Policy>
</PutBucketPolicyRequest>
```

URI Request Parameters

The request uses the following URI parameters.

[name](#)

Specifies the bucket.

For using this parameter with Amazon S3 on Outposts with the REST API, you must specify the name and the x-amz-outpost-id as well.

For using this parameter with S3 on Outposts with the AWS SDK and CLI, you must specify the ARN of the bucket accessed in the format `arn:aws:s3-outposts:<Region>:<account-id>:outpost/<outpost-id>/bucket/<my-bucket-name>`. For example, to access the bucket `reports` through Outpost `my-outpost` owned by account `123456789012` in Region `us-west-2`, use the URL encoding of `arn:aws:s3-outposts:us-west-2:123456789012:outpost/my-outpost/bucket/reports`. The value must be URL encoded.

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

[x-amz-account-id](#)

The AWS account ID of the Outposts bucket.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

x-amz-confirm-remove-self-bucket-access

Set this parameter to true to confirm that you want to remove your permissions to change this bucket policy in the future.

Note

This is not supported by Amazon S3 on Outposts buckets.

Request Body

The request accepts the following data in XML format.

PutBucketPolicyRequest

Root level tag for the PutBucketPolicyRequest parameters.

Required: Yes

Policy

The bucket policy as a JSON document.

Type: String

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Sample request for putting a bucket policy in an Amazon S3 on Outposts bucket

The following request shows the PUT an individual policy request for the Outposts bucket `example-outpost-bucket`.

```
PUT v20180820/bucket/example-outpost-bucket/policy HTTP/1.1
Host: s3-outposts.<Region>.amazonaws.com
Date: Tue, 04 Apr 2010 20:34:56 GMT
Authorization: authorization string
x-amz-account-id: example-account-id
x-amz-outpost-id: op-01ac5d28a6a232904
{
  "Version":"2012-10-17",
  "Id":"exampleS3OnOutpostBucketPolicy",
  "Statement":[
    {
      "Sid":"st1",
      "Effect":"Allow",
      "Principal":{"
        "AWS":"example-account-id"
      },
      "Action":"s3-outposts:*",
      "Resource":"arn:aws:s3-outposts:<your-region>:example-account-id:outpost/
op-01ac5d28a6a232904/bucket/example-outpost-bucket"
    }
  ]
}
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)

- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutBucketReplication

Service: Amazon S3 Control

Note

This action creates an Amazon S3 on Outposts bucket's replication configuration. To create an S3 bucket's replication configuration, see [PutBucketReplication](#) in the *Amazon S3 API Reference*.

Creates a replication configuration or replaces an existing one. For information about S3 replication on Outposts configuration, see [Replicating objects for S3 on Outposts](#) in the *Amazon S3 User Guide*.

Note

It can take a while to propagate PUT or DELETE requests for a replication configuration to all S3 on Outposts systems. Therefore, the replication configuration that's returned by a GET request soon after a PUT or DELETE request might return a more recent result than what's on the Outpost. If an Outpost is offline, the delay in updating the replication configuration on that Outpost can be significant.

Specify the replication configuration in the request body. In the replication configuration, you provide the following information:

- The name of the destination bucket or buckets where you want S3 on Outposts to replicate objects
- The AWS Identity and Access Management (IAM) role that S3 on Outposts can assume to replicate objects on your behalf
- Other relevant information, such as replication rules

A replication configuration must include at least one rule and can contain a maximum of 100. Each rule identifies a subset of objects to replicate by filtering the objects in the source Outposts bucket. To choose additional subsets of objects to replicate, add a rule for each subset.

To specify a subset of the objects in the source Outposts bucket to apply a replication rule to, add the `Filter` element as a child of the `Rule` element. You can filter objects based on an object key prefix, one or more object tags, or both. When you add the `Filter` element in the configuration, you must also add the following elements: `DeleteMarkerReplication`, `Status`, and `Priority`.

Using `PutBucketReplication` on Outposts requires that both the source and destination buckets must have versioning enabled. For information about enabling versioning on a bucket, see [Managing S3 Versioning for your S3 on Outposts bucket](#).

For information about S3 on Outposts replication failure reasons, see [Replication failure reasons](#) in the *Amazon S3 User Guide*.

Handling Replication of Encrypted Objects

Outposts buckets are encrypted at all times. All the objects in the source Outposts bucket are encrypted and can be replicated. Also, all the replicas in the destination Outposts bucket are encrypted with the same encryption key as the objects in the source Outposts bucket.

Permissions

To create a `PutBucketReplication` request, you must have `s3-outposts:PutReplicationConfiguration` permissions for the bucket. The Outposts bucket owner has this permission by default and can grant it to others. For more information about permissions, see [Setting up IAM with S3 on Outposts](#) and [Managing access to S3 on Outposts buckets](#).

Note

To perform this operation, the user or role must also have the `iam:CreateRole` and `iam:PassRole` permissions. For more information, see [Granting a user permissions to pass a role to an AWS service](#).

All Amazon S3 on Outposts REST API requests for this action require an additional parameter of `x-amz-outpost-id` to be passed with the request. In addition, you must use an S3 on Outposts endpoint hostname prefix instead of `s3-control`. For an example of the request syntax for Amazon S3 on Outposts that uses the S3 on Outposts endpoint hostname prefix and the `x-amz-outpost-id` derived by using the access point ARN, see the [Examples](#) section.

The following operations are related to `PutBucketReplication`:

- [GetBucketReplication](#)
- [DeleteBucketReplication](#)

Request Syntax

```
PUT /v20180820/bucket/name/replication HTTP/1.1
Host: Bucket.s3-control.amazonaws.com
x-amz-account-id: AccountId
<?xml version="1.0" encoding="UTF-8"?>
<ReplicationConfiguration xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <Role>string</Role>
  <Rules>
    <Rule>
      <Bucket>string</Bucket>
      <DeleteMarkerReplication>
        <Status>string</Status>
      </DeleteMarkerReplication>
      <Destination>
        <AccessControlTranslation>
          <Owner>string</Owner>
        </AccessControlTranslation>
        <Account>string</Account>
        <Bucket>string</Bucket>
        <EncryptionConfiguration>
          <ReplicaKmsKeyID>string</ReplicaKmsKeyID>
        </EncryptionConfiguration>
        <Metrics>
          <EventThreshold>
            <Minutes>integer</Minutes>
          </EventThreshold>
          <Status>string</Status>
        </Metrics>
        <ReplicationTime>
          <Status>string</Status>
          <Time>
            <Minutes>integer</Minutes>
          </Time>
        </ReplicationTime>
        <StorageClass>string</StorageClass>
      </Destination>
      <ExistingObjectReplication>
        <Status>string</Status>
    </Rule>
  </Rules>
</ReplicationConfiguration>
```

```

</ExistingObjectReplication>
<Filter>
  <And>
    <Prefix>string</Prefix>
    <Tags>
      <S3Tag>
        <Key>string</Key>
        <Value>string</Value>
      </S3Tag>
    </Tags>
  </And>
  <Prefix>string</Prefix>
  <Tag>
    <Key>string</Key>
    <Value>string</Value>
  </Tag>
</Filter>
<ID>string</ID>
<Prefix>string</Prefix>
<Priority>integer</Priority>
<SourceSelectionCriteria>
  <ReplicaModifications>
    <Status>string</Status>
  </ReplicaModifications>
  <SseKmsEncryptedObjects>
    <Status>string</Status>
  </SseKmsEncryptedObjects>
</SourceSelectionCriteria>
<Status>string</Status>
</Rule>
</Rules>
</ReplicationConfiguration>

```

URI Request Parameters

The request uses the following URI parameters.

name

Specifies the S3 on Outposts bucket to set the configuration for.

For using this parameter with Amazon S3 on Outposts with the REST API, you must specify the name and the x-amz-outpost-id as well.

For using this parameter with S3 on Outposts with the AWS SDK and CLI, you must specify the ARN of the bucket accessed in the format `arn:aws:s3-outposts:<Region>:<account-id>:outpost/<outpost-id>/bucket/<my-bucket-name>`. For example, to access the bucket `reports` through Outpost `my-outpost` owned by account `123456789012` in Region `us-west-2`, use the URL encoding of `arn:aws:s3-outposts:us-west-2:123456789012:outpost/my-outpost/bucket/reports`. The value must be URL encoded.

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

[x-amz-account-id](#)

The AWS account ID of the Outposts bucket.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request accepts the following data in XML format.

[ReplicationConfiguration](#)

Root level tag for the `ReplicationConfiguration` parameters.

Required: Yes

[Role](#)

The Amazon Resource Name (ARN) of the AWS Identity and Access Management (IAM) role that S3 on Outposts assumes when replicating objects. For information about S3 replication on Outposts configuration, see [Setting up replication](#) in the *Amazon S3 User Guide*.

Type: String

Required: Yes

Rules

A container for one or more replication rules. A replication configuration must have at least one rule and can contain an array of 100 rules at the most.

Type: Array of [ReplicationRule](#) data types

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Sample Request: Add a replication configuration to an Amazon S3 on Outposts bucket

The following sample PUT request creates a replication subresource on the specified Outposts bucket named `example-outpost-bucket` and saves the replication configuration in it. The replication configuration specifies a rule to replicate objects to the `example-outpost-bucket` bucket. The rule includes a filter to replicate only the objects that are created with the key name prefix `TaxDocs` and that have two specific tags.

After you add a replication configuration to your Outposts bucket, S3 on Outposts assumes the AWS Identity and Access Management (IAM) role that's specified in the configuration to replicate objects on behalf of the Outposts bucket owner. The bucket owner is the AWS account that created the Outposts bucket.

Filtering by using the `Filter` element is supported in the latest XML configuration. The earlier version of the XML configuration isn't supported.

For more examples of S3 replication on Outposts configuration, see [Creating replication rules on Outposts](#) in the *Amazon S3 User Guide*.

```
PUT /v20180820/bucket/example-outpost-bucket/replication HTTP/1.1
Host:s3-outposts.<Region>.amazonaws.com
```

```

x-amz-account-id: example-account-id
x-amz-outpost-id: op-01ac5d28a6a232904
Authorization: authorization string

<ReplicationConfiguration>
  <Role>arn:aws:iam::35667example:role/ReplicationRoleForS3Outposts</Role>
  <Rules>
    <Rule>
      <Bucket>arn:aws:s3-outposts:us-east-1:example-account-id:outpost/SOURCE-OUTPOST-
ID/accesspoint/SOURCE-OUTPOSTS-BUCKET-ACCESS-POINT</Bucket>
      <ID>rule1</ID>
      <Status>Enabled</Status>
      <Priority>1</Priority>
      <DeleteMarkerReplication>
        <Status>Disabled</Status>
      </DeleteMarkerReplication>
      <Filter>
        <And>
          <Prefix>TaxDocs</Prefix>
          <Tag>
            <Key>key1</Key>
            <Value>value1</Value>
          </Tag>
          <Tag>
            <Key>key2</Key>
            <Value>value2</Value>
          </Tag>
        </And>
      </Filter>
      <Destination>
        <Bucket>arn:aws:s3-outposts:us-east-1:example-account-id:outpost/DESTINATION-
OUTPOST-ID/accesspoint/DESTINATION-OUTPOSTS-BUCKET-ACCESS-POINT</Bucket>
      </Destination>
    </Rule>
  </Rules>
</ReplicationConfiguration>

```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutBucketTagging

Service: Amazon S3 Control

Note

This action puts tags on an Amazon S3 on Outposts bucket. To put tags on an S3 bucket, see [PutBucketTagging](#) in the *Amazon S3 API Reference*.

Sets the tags for an S3 on Outposts bucket. For more information, see [Using Amazon S3 on Outposts](#) in the *Amazon S3 User Guide*.

Use tags to organize your AWS bill to reflect your own cost structure. To do this, sign up to get your AWS account bill with tag key values included. Then, to see the cost of combined resources, organize your billing information according to resources with the same tag key values. For example, you can tag several resources with a specific application name, and then organize your billing information to see the total cost of that application across several services. For more information, see [Cost allocation and tagging](#).

Note

Within a bucket, if you add a tag that has the same key as an existing tag, the new value overwrites the old value. For more information, see [Using cost allocation in Amazon S3 bucket tags](#).

To use this action, you must have permissions to perform the `s3-outposts:PutBucketTagging` action. The Outposts bucket owner has this permission by default and can grant this permission to others. For more information about permissions, see [Permissions Related to Bucket Subresource Operations](#) and [Managing access permissions to your Amazon S3 resources](#).

PutBucketTagging has the following special errors:

- Error code: `InvalidTagError`
 - Description: The tag provided was not a valid tag. This error can occur if the tag did not pass input validation. For information about tag restrictions, see [User-Defined Tag Restrictions](#) and [AWS-Generated Cost Allocation Tag Restrictions](#).
- Error code: `MalformedXMLError`

- Description: The XML provided does not match the schema.
- Error code: `OperationAbortedError`
 - Description: A conflicting conditional action is currently in progress against this resource. Try again.
- Error code: `InternalError`
 - Description: The service was unable to apply the provided tag to the bucket.

All Amazon S3 on Outposts REST API requests for this action require an additional parameter of `x-amz-outpost-id` to be passed with the request. In addition, you must use an S3 on Outposts endpoint hostname prefix instead of `s3-control`. For an example of the request syntax for Amazon S3 on Outposts that uses the S3 on Outposts endpoint hostname prefix and the `x-amz-outpost-id` derived by using the access point ARN, see the [Examples](#) section.

The following actions are related to `PutBucketTagging`:

- [GetBucketTagging](#)
- [DeleteBucketTagging](#)

Request Syntax

```
PUT /v20180820/bucket/name/tagging HTTP/1.1
Host: Bucket.s3-control.amazonaws.com
x-amz-account-id: AccountId
<?xml version="1.0" encoding="UTF-8"?>
<Tagging xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <TagSet>
    <S3Tag>
      <Key>string</Key>
      <Value>string</Value>
    </S3Tag>
  </TagSet>
</Tagging>
```

URI Request Parameters

The request uses the following URI parameters.

name

The Amazon Resource Name (ARN) of the bucket.

For using this parameter with Amazon S3 on Outposts with the REST API, you must specify the name and the x-amz-outpost-id as well.

For using this parameter with S3 on Outposts with the AWS SDK and CLI, you must specify the ARN of the bucket accessed in the format `arn:aws:s3-outposts:<Region>:<account-id>:outpost/<outpost-id>/bucket/<my-bucket-name>`. For example, to access the bucket `reports` through Outpost `my-outpost` owned by account `123456789012` in Region `us-west-2`, use the URL encoding of `arn:aws:s3-outposts:us-west-2:123456789012:outpost/my-outpost/bucket/reports`. The value must be URL encoded.

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

x-amz-account-id

The AWS account ID of the Outposts bucket.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request accepts the following data in XML format.

Tagging

Root level tag for the Tagging parameters.

Required: Yes

TagSet

A collection for a set of tags.

Type: Array of [S3Tag](#) data types

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Sample request: Add tag set to an Amazon S3 on Outposts bucket

The following request adds a tag set to the existing `example-outpost-bucket` bucket.

```
PUT v20180820/bucket/example-outpost-bucket/tagging HTTP/1.1
Host: s3-outposts.<Region>.amazonaws.com
Content-Length: 1660
x-amz-date: Thu, 12 Nov 2020 20:04:21 GMT
x-amz-account-id: example-account-id
x-amz-outpost-id: op-01ac5d28a6a232904
Authorization: authorization string

<Tagging>
  <TagSet>
    <Tag>
      <Key>Project</Key>
      <Value>Project One</Value>
    </Tag>
    <Tag>
      <Key>User</Key>
      <Value>jsmith</Value>
    </Tag>
  </TagSet>
</Tagging>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutBucketVersioning

Service: Amazon S3 Control

Note

This operation sets the versioning state for S3 on Outposts buckets only. To set the versioning state for an S3 bucket, see [PutBucketVersioning](#) in the *Amazon S3 API Reference*.

Sets the versioning state for an S3 on Outposts bucket. With S3 Versioning, you can save multiple distinct copies of your objects and recover from unintended user actions and application failures.

You can set the versioning state to one of the following:

- **Enabled** - Enables versioning for the objects in the bucket. All objects added to the bucket receive a unique version ID.
- **Suspended** - Suspends versioning for the objects in the bucket. All objects added to the bucket receive the version ID `null`.

If you've never set versioning on your bucket, it has no versioning state. In that case, a [GetBucketVersioning](#) request does not return a versioning state value.

When you enable S3 Versioning, for each object in your bucket, you have a current version and zero or more noncurrent versions. You can configure your bucket S3 Lifecycle rules to expire noncurrent versions after a specified time period. For more information, see [Creating and managing a lifecycle configuration for your S3 on Outposts bucket](#) in the *Amazon S3 User Guide*.

If you have an object expiration lifecycle configuration in your non-versioned bucket and you want to maintain the same permanent delete behavior when you enable versioning, you must add a noncurrent expiration policy. The noncurrent expiration lifecycle configuration will manage the deletes of the noncurrent object versions in the version-enabled bucket. For more information, see [Versioning](#) in the *Amazon S3 User Guide*.

All Amazon S3 on Outposts REST API requests for this action require an additional parameter of `x-amz-outpost-id` to be passed with the request. In addition, you must use an S3 on Outposts endpoint hostname prefix instead of `s3-control`. For an example of the request syntax for Amazon S3 on Outposts that uses the S3 on Outposts endpoint hostname prefix and the `x-amz-outpost-id` derived by using the access point ARN, see the [Examples](#) section.

The following operations are related to PutBucketVersioning for S3 on Outposts.

- [GetBucketVersioning](#)
- [PutBucketLifecycleConfiguration](#)
- [GetBucketLifecycleConfiguration](#)

Request Syntax

```
PUT /v20180820/bucket/name/versioning HTTP/1.1
Host: Bucket.s3-control.amazonaws.com
x-amz-account-id: AccountId
x-amz-mfa: MFA
<?xml version="1.0" encoding="UTF-8"?>
<VersioningConfiguration xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <MfaDelete>string</MfaDelete>
  <Status>string</Status>
</VersioningConfiguration>
```

URI Request Parameters

The request uses the following URI parameters.

[name](#)

The S3 on Outposts bucket to set the versioning state for.

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

[x-amz-account-id](#)

The AWS account ID of the S3 on Outposts bucket.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

x-amz-mfa

The concatenation of the authentication device's serial number, a space, and the value that is displayed on your authentication device.

Request Body

The request accepts the following data in XML format.

VersioningConfiguration

Root level tag for the VersioningConfiguration parameters.

Required: Yes

MFADelete

Specifies whether MFA delete is enabled or disabled in the bucket versioning configuration for the S3 on Outposts bucket.

Type: String

Valid Values: Enabled | Disabled

Required: No

Status

Sets the versioning state of the S3 on Outposts bucket.

Type: String

Valid Values: Enabled | Suspended

Required: No

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Sample PutBucketVersioning request on an Amazon S3 on Outposts bucket

This request sets the versioning state on an S3 on Outposts bucket that's named `example-outpost-bucket`.

```
PUT /v20180820/bucket/example-outpost-bucket/?versioning HTTP/1.1
Host:s3-outposts.region-code.amazonaws.com
x-amz-account-id: example-account-id
x-amz-outpost-id: op-01ac5d28a6a232904
Content-Length: 0
Date: Wed, 25 May 2022 12:00:00 GMT
Content-MD5: q6yJDlIkBaGGfb3QLY69A==
Authorization: authorization string
Content-Length: 214

<VersioningConfiguration xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Status>Enabled</Status>
</VersioningConfiguration>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutJobTagging

Service: Amazon S3 Control

Sets the supplied tag-set on an S3 Batch Operations job.

A tag is a key-value pair. You can associate S3 Batch Operations tags with any job by sending a PUT request against the tagging subresource that is associated with the job. To modify the existing tag set, you can either replace the existing tag set entirely, or make changes within the existing tag set by retrieving the existing tag set using [GetJobTagging](#), modify that tag set, and use this operation to replace the tag set with the one you modified. For more information, see [Controlling access and labeling jobs using tags](#) in the *Amazon S3 User Guide*.

Note

- If you send this request with an empty tag set, Amazon S3 deletes the existing tag set on the Batch Operations job. If you use this method, you are charged for a Tier 1 Request (PUT). For more information, see [Amazon S3 pricing](#).
- For deleting existing tags for your Batch Operations job, a [DeleteJobTagging](#) request is preferred because it achieves the same result without incurring charges.
- A few things to consider about using tags:
 - Amazon S3 limits the maximum number of tags to 50 tags per job.
 - You can associate up to 50 tags with a job as long as they have unique tag keys.
 - A tag key can be up to 128 Unicode characters in length, and tag values can be up to 256 Unicode characters in length.
 - The key and values are case sensitive.
 - For tagging-related restrictions related to characters and encodings, see [User-Defined Tag Restrictions](#) in the *AWS Billing and Cost Management User Guide*.

Permissions

To use the PutJobTagging operation, you must have permission to perform the `s3:PutJobTagging` action.

Related actions include:

- [CreateJob](#)
- [GetJobTagging](#)
- [DeleteJobTagging](#)

Request Syntax

```
PUT /v20180820/jobs/id/tagging HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
<?xml version="1.0" encoding="UTF-8"?>
<PutJobTaggingRequest xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <Tags>
    <S3Tag>
      <Key>string</Key>
      <Value>string</Value>
    </S3Tag>
  </Tags>
</PutJobTaggingRequest>
```

URI Request Parameters

The request uses the following URI parameters.

[id](#)

The ID for the S3 Batch Operations job whose tags you want to replace.

Length Constraints: Minimum length of 5. Maximum length of 36.

Pattern: `[a-zA-Z0-9\-\_\+]`

Required: Yes

[x-amz-account-id](#)

The AWS account ID associated with the S3 Batch Operations job.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request accepts the following data in XML format.

[PutJobTaggingRequest](#)

Root level tag for the PutJobTaggingRequest parameters.

Required: Yes

[Tags](#)

The set of tags to associate with the S3 Batch Operations job.

Type: Array of [S3Tag](#) data types

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

InternalServerErrorException

HTTP Status Code: 500

NotFoundException

HTTP Status Code: 400

TooManyRequestsException

HTTP Status Code: 400

TooManyTagsException

Amazon S3 throws this exception if you have too many tags in your tag set.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutMultiRegionAccessPointPolicy

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Associates an access control policy with the specified Multi-Region Access Point. Each Multi-Region Access Point can have only one policy, so a request made to this action replaces any existing policy that is associated with the specified Multi-Region Access Point.

This action will always be routed to the US West (Oregon) Region. For more information about the restrictions around working with Multi-Region Access Points, see [Multi-Region Access Point restrictions and limitations](#) in the *Amazon S3 User Guide*.

The following actions are related to PutMultiRegionAccessPointPolicy:

- [GetMultiRegionAccessPointPolicy](#)
- [GetMultiRegionAccessPointPolicyStatus](#)

Request Syntax

```
POST /v20180820/async-requests/mrap/put-policy HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
<?xml version="1.0" encoding="UTF-8"?>
<PutMultiRegionAccessPointPolicyRequest xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <ClientToken>string</ClientToken>
  <Details>
    <Name>string</Name>
    <Policy>string</Policy>
  </Details>
</PutMultiRegionAccessPointPolicyRequest>
```

URI Request Parameters

The request uses the following URI parameters.

x-amz-account-id

The AWS account ID for the owner of the Multi-Region Access Point.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request accepts the following data in XML format.

PutMultiRegionAccessPointPolicyRequest

Root level tag for the PutMultiRegionAccessPointPolicyRequest parameters.

Required: Yes

ClientToken

An idempotency token used to identify the request and guarantee that requests are unique.

Type: String

Length Constraints: Maximum length of 64.

Pattern: `\S+`

Required: Yes

Details

A container element containing the details of the policy for the Multi-Region Access Point.

Type: [PutMultiRegionAccessPointPolicyInput](#) data type

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

```
<?xml version="1.0" encoding="UTF-8"?>
<PutMultiRegionAccessPointPolicyResult>
  <RequestTokenARN>string</RequestTokenARN>
</PutMultiRegionAccessPointPolicyResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

PutMultiRegionAccessPointPolicyResult

Root level tag for the PutMultiRegionAccessPointPolicyResult parameters.

Required: Yes

RequestTokenARN

The request token associated with the request. You can use this token with [DescribeMultiRegionAccessPointOperation](#) to determine the status of asynchronous requests.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: arn: .+

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)

- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutPublicAccessBlock

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Creates or modifies the `PublicAccessBlock` configuration for an AWS account. This operation may be restricted when the account is managed by organization-level Block Public Access policies. You might get an Access Denied (403) error when the account is managed by organization-level Block Public Access policies. Organization-level policies override account-level settings, preventing direct account-level modifications. For this operation, users must have the `s3:PutAccountPublicAccessBlock` permission. For more information, see [Using Amazon S3 block public access](#).

Related actions include:

- [GetPublicAccessBlock](#)
- [DeletePublicAccessBlock](#)

Request Syntax

```
PUT /v20180820/configuration/publicAccessBlock HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
<?xml version="1.0" encoding="UTF-8"?>
<PublicAccessBlockConfiguration xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <BlockPublicAcls>boolean</BlockPublicAcls>
  <IgnorePublicAcls>boolean</IgnorePublicAcls>
  <BlockPublicPolicy>boolean</BlockPublicPolicy>
  <RestrictPublicBuckets>boolean</RestrictPublicBuckets>
</PublicAccessBlockConfiguration>
```

URI Request Parameters

The request uses the following URI parameters.

x-amz-account-id

The account ID for the AWS account whose PublicAccessBlock configuration you want to set.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request accepts the following data in XML format.

PublicAccessBlockConfiguration

Root level tag for the PublicAccessBlockConfiguration parameters.

Required: Yes

BlockPublicAcls

Specifies whether Amazon S3 should block public access control lists (ACLs) for buckets in this account. Setting this element to TRUE causes the following behavior:

- PutBucketAcl and PutObjectAcl calls fail if the specified ACL is public.
- PUT Object calls fail if the request includes a public ACL.
- PUT Bucket calls fail if the request includes a public ACL.

Enabling this setting doesn't affect existing policies or ACLs.

This property is not supported for Amazon S3 on Outposts.

Type: Boolean

Required: No

BlockPublicPolicy

Specifies whether Amazon S3 should block public bucket policies for buckets in this account. Setting this element to TRUE causes Amazon S3 to reject calls to PUT Bucket policy if the specified bucket policy allows public access.

Enabling this setting doesn't affect existing bucket policies.

This property is not supported for Amazon S3 on Outposts.

Type: Boolean

Required: No

IgnorePublicAcls

Specifies whether Amazon S3 should ignore public ACLs for buckets in this account. Setting this element to TRUE causes Amazon S3 to ignore all public ACLs on buckets in this account and any objects that they contain.

Enabling this setting doesn't affect the persistence of any existing ACLs and doesn't prevent new public ACLs from being set.

This property is not supported for Amazon S3 on Outposts.

Type: Boolean

Required: No

RestrictPublicBuckets

Specifies whether Amazon S3 should restrict public bucket policies for buckets in this account. Setting this element to TRUE restricts access to buckets with public policies to only AWS service principals and authorized users within this account.

Enabling this setting doesn't affect previously stored bucket policies, except that public and cross-account access within any public bucket policy, including non-public delegation to specific accounts, is blocked.

This property is not supported for Amazon S3 on Outposts.

Type: Boolean

Required: No

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutStorageLensConfiguration

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Puts an Amazon S3 Storage Lens configuration. For more information about S3 Storage Lens, see [Working with Amazon S3 Storage Lens](#) in the *Amazon S3 User Guide*. For a complete list of S3 Storage Lens metrics, see [S3 Storage Lens metrics glossary](#) in the *Amazon S3 User Guide*.

Note

To use this action, you must have permission to perform the `s3:PutStorageLensConfiguration` action. For more information, see [Setting permissions to use Amazon S3 Storage Lens](#) in the *Amazon S3 User Guide*.

Request Syntax

```
PUT /v20180820/storagelens/storagelensid HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
<?xml version="1.0" encoding="UTF-8"?>
<PutStorageLensConfigurationRequest xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <StorageLensConfiguration>
    <AccountLevel>
      <ActivityMetrics>
        <IsEnabled>boolean</IsEnabled>
      </ActivityMetrics>
      <AdvancedCostOptimizationMetrics>
        <IsEnabled>boolean</IsEnabled>
      </AdvancedCostOptimizationMetrics>
      <AdvancedDataProtectionMetrics>
        <IsEnabled>boolean</IsEnabled>
      </AdvancedDataProtectionMetrics>
      <AdvancedPerformanceMetrics>
        <IsEnabled>boolean</IsEnabled>
      </AdvancedPerformanceMetrics>
    </AccountLevel>
  </StorageLensConfiguration>
</PutStorageLensConfigurationRequest>
```

```

<BucketLevel>
  <ActivityMetrics>
    <IsEnabled>boolean</IsEnabled>
  </ActivityMetrics>
  <AdvancedCostOptimizationMetrics>
    <IsEnabled>boolean</IsEnabled>
  </AdvancedCostOptimizationMetrics>
  <AdvancedDataProtectionMetrics>
    <IsEnabled>boolean</IsEnabled>
  </AdvancedDataProtectionMetrics>
  <AdvancedPerformanceMetrics>
    <IsEnabled>boolean</IsEnabled>
  </AdvancedPerformanceMetrics>
  <DetailedStatusCodesMetrics>
    <IsEnabled>boolean</IsEnabled>
  </DetailedStatusCodesMetrics>
  <PrefixLevel>
    <StorageMetrics>
      <IsEnabled>boolean</IsEnabled>
      <SelectionCriteria>
        <Delimiter>string</Delimiter>
        <MaxDepth>integer</MaxDepth>
        <MinStorageBytesPercentage>double</MinStorageBytesPercentage>
      </SelectionCriteria>
    </StorageMetrics>
  </PrefixLevel>
</BucketLevel>
<DetailedStatusCodesMetrics>
  <IsEnabled>boolean</IsEnabled>
</DetailedStatusCodesMetrics>
<StorageLensGroupLevel>
  <SelectionCriteria>
    <Exclude>
      <Arn>string</Arn>
    </Exclude>
    <Include>
      <Arn>string</Arn>
    </Include>
  </SelectionCriteria>
</StorageLensGroupLevel>
</AccountLevel>
<AwsOrg>
  <Arn>string</Arn>
</AwsOrg>

```

```

<DataExport>
  <CloudWatchMetrics>
    <IsEnabled>boolean</IsEnabled>
  </CloudWatchMetrics>
  <S3BucketDestination>
    <AccountId>string</AccountId>
    <Arn>string</Arn>
    <Encryption>
      <SSE-KMS>
        <KeyId>string</KeyId>
      </SSE-KMS>
      <SSE-S3>
      </SSE-S3>
    </Encryption>
    <Format>string</Format>
    <OutputSchemaVersion>string</OutputSchemaVersion>
    <Prefix>string</Prefix>
  </S3BucketDestination>
  <StorageLensTableDestination>
    <Encryption>
      <SSE-KMS>
        <KeyId>string</KeyId>
      </SSE-KMS>
      <SSE-S3>
      </SSE-S3>
    </Encryption>
    <IsEnabled>boolean</IsEnabled>
  </StorageLensTableDestination>
</DataExport>
<Exclude>
  <Buckets>
    <Arn>string</Arn>
  </Buckets>
  <Regions>
    <Region>string</Region>
  </Regions>
</Exclude>
<ExpandedPrefixesDataExport>
  <S3BucketDestination>
    <AccountId>string</AccountId>
    <Arn>string</Arn>
    <Encryption>
      <SSE-KMS>
        <KeyId>string</KeyId>

```

```

        </SSE-KMS>
        <SSE-S3>
        </SSE-S3>
    </Encryption>
    <Format>string</Format>
    <OutputSchemaVersion>string</OutputSchemaVersion>
    <Prefix>string</Prefix>
</S3BucketDestination>
<StorageLensTableDestination>
    <Encryption>
        <SSE-KMS>
            <KeyId>string</KeyId>
        </SSE-KMS>
        <SSE-S3>
        </SSE-S3>
    </Encryption>
    <IsEnabled>boolean</IsEnabled>
</StorageLensTableDestination>
</ExpandedPrefixesDataExport>
<Id>string</Id>
<Include>
    <Buckets>
        <Arn>string</Arn>
    </Buckets>
    <Regions>
        <Region>string</Region>
    </Regions>
</Include>
<IsEnabled>boolean</IsEnabled>
<PrefixDelimiter>string</PrefixDelimiter>
<StorageLensArn>string</StorageLensArn>
</StorageLensConfiguration>
<Tags>
    <Tag>
        <Key>string</Key>
        <Value>string</Value>
    </Tag>
</Tags>
</PutStorageLensConfigurationRequest>

```

URI Request Parameters

The request uses the following URI parameters.

storageensid

The ID of the S3 Storage Lens configuration.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [a-zA-Z0-9\-_\.\.]+

Required: Yes

x-amz-account-id

The account ID of the requester.

Length Constraints: Maximum length of 64.

Pattern: ^\d{12}\$

Required: Yes

Request Body

The request accepts the following data in XML format.

PutStorageLensConfigurationRequest

Root level tag for the PutStorageLensConfigurationRequest parameters.

Required: Yes

StorageLensConfiguration

The S3 Storage Lens configuration.

Type: [StorageLensConfiguration](#) data type

Required: Yes

Tags

The tag set of the S3 Storage Lens configuration.

Note

You can set up to a maximum of 50 tags.

Type: Array of [StorageLensTag](#) data types

Required: No

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutStorageLensConfigurationTagging

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Put or replace tags on an existing Amazon S3 Storage Lens configuration. For more information about S3 Storage Lens, see [Assessing your storage activity and usage with Amazon S3 Storage Lens](#) in the *Amazon S3 User Guide*.

Note

To use this action, you must have permission to perform the `s3:PutStorageLensConfigurationTagging` action. For more information, see [Setting permissions to use Amazon S3 Storage Lens](#) in the *Amazon S3 User Guide*.

Request Syntax

```
PUT /v20180820/storageLens/storageLensId/tagging HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
<?xml version="1.0" encoding="UTF-8"?>
<PutStorageLensConfigurationTaggingRequest xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <Tags>
    <Tag>
      <Key>string</Key>
      <Value>string</Value>
    </Tag>
  </Tags>
</PutStorageLensConfigurationTaggingRequest>
```

URI Request Parameters

The request uses the following URI parameters.

storagelensid

The ID of the S3 Storage Lens configuration.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\_\.\ ]+`

Required: Yes

x-amz-account-id

The account ID of the requester.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request accepts the following data in XML format.

PutStorageLensConfigurationTaggingRequest

Root level tag for the PutStorageLensConfigurationTaggingRequest parameters.

Required: Yes

Tags

The tag set of the S3 Storage Lens configuration.

Note

You can set up to a maximum of 50 tags.

Type: Array of [StorageLensTag](#) data types

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

SubmitMultiRegionAccessPointRoutes

Service: Amazon S3 Control

Note

This operation is not supported by directory buckets.

Submits an updated route configuration for a Multi-Region Access Point. This API operation updates the routing status for the specified Regions from active to passive, or from passive to active. A value of 0 indicates a passive status, which means that traffic won't be routed to the specified Region. A value of 100 indicates an active status, which means that traffic will be routed to the specified Region. At least one Region must be active at all times.

When the routing configuration is changed, any in-progress operations (uploads, copies, deletes, and so on) to formerly active Regions will continue to run to their final completion state (success or failure). The routing configurations of any Regions that aren't specified remain unchanged.

Note

Updated routing configurations might not be immediately applied. It can take up to 2 minutes for your changes to take effect.

To submit routing control changes and failover requests, use the Amazon S3 failover control infrastructure endpoints in these five AWS Regions:

- us-east-1
- us-west-2
- ap-southeast-2
- ap-northeast-1
- eu-west-1

Request Syntax

```
PATCH /v20180820/mrap/instances/mrap+/routes HTTP/1.1
Host: s3-control.amazonaws.com
```

```
x-amz-account-id: AccountId
<?xml version="1.0" encoding="UTF-8"?>
<SubmitMultiRegionAccessPointRoutesRequest xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <RouteUpdates>
    <Route>
      <Bucket>string</Bucket>
      <Region>string</Region>
      <TrafficDialPercentage>integer</TrafficDialPercentage>
    </Route>
  </RouteUpdates>
</SubmitMultiRegionAccessPointRoutesRequest>
```

URI Request Parameters

The request uses the following URI parameters.

mrp

The Multi-Region Access Point ARN.

Length Constraints: Maximum length of 200.

Pattern: `^[a-zA-Z0-9\:\.-]{3,200}$`

Required: Yes

x-amz-account-id

The AWS account ID for the owner of the Multi-Region Access Point.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request accepts the following data in XML format.

SubmitMultiRegionAccessPointRoutesRequest

Root level tag for the SubmitMultiRegionAccessPointRoutesRequest parameters.

Required: Yes

RouteUpdates

The different routes that make up the new route configuration. Active routes return a value of 100, and passive routes return a value of 0.

Type: Array of [MultiRegionAccessPointRoute](#) data types

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Examples

Sample request for initiating failover

In the following example, the request to submit these routing changes to initiate a failover is sent to the failover control infrastructure in the us-east-1 Region. In this example, the eu-north-1 Region is set to active, and the ap-northeast-3 Region is set to passive. In other words, the ap-northeast-3 Region is failed over to the eu-north-1 Region.

```
PATCH /v20180820/mrap/instances/<Multi-Region Access Point>/routes HTTP/1.1
Host: example-account-id.s3-control.us-east-1.amazonaws.com
```

```
<SubmitMultiRegionAccessPointRoutesRequest>
  <RouteUpdates>
    <Route>
      <Region>eu-north-1</Region>
      <Bucket>example-bucket-eu-north-1</Bucket>
      <TrafficDialPercentage>100</TrafficDialPercentage>
    </Route>
    <Route>
      <Region>ap-northeast-3</Region>
      <Bucket>example-bucket-ap-northeast-3</Bucket>
```



```
<TrafficDialPercentage>0</TrafficDialPercentage>
</Route>
</RouteUpdates>
</SubmitMultiRegionAccessPointRoutesRequest>
```

Sample request for setting a Region to active status

The following request updates the route configuration of the eu-north-1 Region to active. The request is sent to the failover control infrastructure in the eu-west-1 Region.

```
PATCH /v20180820/mrap/instances/<Multi-Region Access Point>/routes HTTP/1.1
Host: example-account-id.s3-control.eu-west-1.amazonaws.com
```

```
<SubmitMultiRegionAccessPointRoutesRequest>
  <RouteUpdates>
    <Route>
      <Region>eu-north-1</Region>
      <Bucket>example-bucket-eu-north-1</Bucket>
      <TrafficDialPercentage>100</TrafficDialPercentage>
    </Route>
  </RouteUpdates>
</SubmitMultiRegionAccessPointRoutesRequest>
```

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)

- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

TagResource

Service: Amazon S3 Control

Creates a new user-defined tag or updates an existing tag. Each tag is a label consisting of a key and value that is applied to your resource. Tags can help you organize, track costs for, and control access to your resources. You can add up to 50 AWS resource tags for each S3 resource.

Note

This operation is only supported for the following Amazon S3 resource:

- [General purpose buckets](#)
- [Access Points for directory buckets](#)
- [Access Points for general purpose buckets](#)
- [Directory buckets](#)
- [S3 Storage Lens groups](#)
- [S3 Access Grants instances, registered locations, or grants.](#)

Permissions

For general purpose buckets, access points for general purpose buckets, Storage Lens groups, and S3 Access Grants, you must have the `s3:TagResource` permission to use this operation.

Directory bucket permissions

For directory buckets, you must have the `s3express:TagResource` permission to use this operation. For more information about directory buckets policies and permissions, see [Identity and Access Management \(IAM\) for S3 Express One Zone](#) in the *Amazon S3 User Guide*.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `s3express-control.region.amazonaws.com`.

For information about S3 Tagging errors, see [List of Amazon S3 Tagging error codes](#).

Request Syntax

```
POST /v20180820/tags/resourceArn HTTP/1.1
```

```
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
<?xml version="1.0" encoding="UTF-8"?>
<TagResourceRequest xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <Tags>
    <Tag>
      <Key>string</Key>
      <Value>string</Value>
    </Tag>
  </Tags>
</TagResourceRequest>
```

URI Request Parameters

The request uses the following URI parameters.

[resourceArn](#)

The Amazon Resource Name (ARN) of the S3 resource that you're applying tags to. The tagged resource can be a directory bucket, S3 Storage Lens group or S3 Access Grants instance, registered location, or grant.

Length Constraints: Maximum length of 1011.

Pattern: `arn:[^:]+:s3(express)?:[^:]*`

Required: Yes

[x-amz-account-id](#)

The AWS account ID that created the S3 resource that you're trying to add tags to or the requester's account ID.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request accepts the following data in XML format.

[TagResourceRequest](#)

Root level tag for the TagResourceRequest parameters.

Required: Yes

[Tags](#)

The AWS resource tags that you want to add to the specified S3 resource.

Type: Array of [Tag](#) data types

Array Members: Minimum number of 0 items. Maximum number of 50 items.

Required: Yes

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)

- [AWS SDK for Ruby V3](#)

UntagResource

Service: Amazon S3 Control

This operation removes the specified user-defined tags from an S3 resource. You can pass one or more tag keys.

Note

This operation is only supported for the following Amazon S3 resources:

- [General purpose buckets](#)
- [Access Points for directory buckets](#)
- [Access Points for general purpose buckets](#)
- [Directory buckets](#)
- [S3 Storage Lens groups](#)
- [S3 Access Grants instances, registered locations, and grants.](#)

Permissions

For general purpose buckets, access points for general purpose buckets, Storage Lens groups, and S3 Access Grants, you must have the `s3:UntagResource` permission to use this operation.

Directory bucket permissions

For directory buckets, you must have the `s3express:UntagResource` permission to use this operation. For more information about directory buckets policies and permissions, see [Identity and Access Management \(IAM\) for S3 Express One Zone](#) in the *Amazon S3 User Guide*.

HTTP Host header syntax

Directory buckets - The HTTP Host header syntax is `s3express-control.region.amazonaws.com`.

For information about S3 Tagging errors, see [List of Amazon S3 Tagging error codes](#).

Request Syntax

```
DELETE /v20180820/tags/resourceArn?tagKeys=TagKeys HTTP/1.1
```

```
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

[resourceArn](#)

The Amazon Resource Name (ARN) of the S3 resource that you're removing tags from. The tagged resource can be a directory bucket, S3 Storage Lens group or S3 Access Grants instance, registered location, or grant.

Length Constraints: Maximum length of 1011.

Pattern: `arn:[^:]+:s3(express)?:[^:]*`

Required: Yes

[tagKeys](#)

The array of tag key-value pairs that you're trying to remove from of the S3 resource.

Array Members: Minimum number of 0 items. Maximum number of 50 items.

Length Constraints: Minimum length of 1. Maximum length of 128.

Pattern: `^([\p{L}\p{Z}\p{N}_.: /+=\ -@]*)$`

Required: Yes

[x-amz-account-id](#)

The AWS account ID that owns the resource that you're trying to remove the tags from.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

UpdateAccessGrantsLocation

Service: Amazon S3 Control

Updates the IAM role of a registered location in your S3 Access Grants instance.

Permissions

You must have the `s3:UpdateAccessGrantsLocation` permission to use this operation.

Additional Permissions

You must also have the following permission: `iam:PassRole`

Request Syntax

```
PUT /v20180820/accessgrantsinstance/location/id HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
<?xml version="1.0" encoding="UTF-8"?>
<UpdateAccessGrantsLocationRequest xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <IAMRoleArn>string</IAMRoleArn>
</UpdateAccessGrantsLocationRequest>
```

URI Request Parameters

The request uses the following URI parameters.

id

The ID of the registered location that you are updating. S3 Access Grants assigns this ID when you register the location. S3 Access Grants assigns the ID default to the default location `s3://` and assigns an auto-generated ID to other locations that you register.

The ID of the registered location to which you are granting access. S3 Access Grants assigned this ID when you registered the location. S3 Access Grants assigns the ID default to the default location `s3://` and assigns an auto-generated ID to other locations that you register.

If you are passing the default location, you cannot create an access grant for the entire default location. You must also specify a bucket or a bucket and prefix in the Subprefix field.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [a-zA-Z0-9\-]+

Required: Yes

x-amz-account-id

The AWS account ID of the S3 Access Grants instance.

Length Constraints: Maximum length of 64.

Pattern: ^\d{12}\$

Required: Yes

Request Body

The request accepts the following data in XML format.

UpdateAccessGrantsLocationRequest

Root level tag for the UpdateAccessGrantsLocationRequest parameters.

Required: Yes

IAMRoleArn

The Amazon Resource Name (ARN) of the IAM role for the registered location. S3 Access Grants assumes this role to manage access to the registered location.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: arn:[^:]+:iam::\d{12}:role/.*

Required: Yes

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<UpdateAccessGrantsLocationResult>
  <CreatedAt>timestamp</CreatedAt>
```

```
<AccessGrantsLocationId>string</AccessGrantsLocationId>  
<AccessGrantsLocationArn>string</AccessGrantsLocationArn>  
<LocationScope>string</LocationScope>  
<IAMRoleArn>string</IAMRoleArn>  
</UpdateAccessGrantsLocationResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

UpdateAccessGrantsLocationResult

Root level tag for the UpdateAccessGrantsLocationResult parameters.

Required: Yes

AccessGrantsLocationArn

The Amazon Resource Name (ARN) of the registered location that you are updating.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:[a-z\-\+]:s3:[a-z0-9\-\+]:\d{12}:access\-\grants\location/[a-zA-Z0-9\-\+]`

AccessGrantsLocationId

The ID of the registered location to which you are granting access. S3 Access Grants assigned this ID when you registered the location. S3 Access Grants assigns the ID default to the default location `s3://` and assigns an auto-generated ID to other locations that you register.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\+]`

CreatedAt

The date and time when you registered the location.

Type: Timestamp

IAMRoleArn

The Amazon Resource Name (ARN) of the IAM role of the registered location. S3 Access Grants assumes this role to manage access to the registered location.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:[^:]+:iam:\d{12}:role/.*`

LocationScope

The S3 URI path of the location that you are updating. You cannot update the scope of the registered location. The location scope can be the default S3 location `s3://`, the S3 path to a bucket `s3://<bucket>`, or the S3 path to a bucket and prefix `s3://<bucket>/<prefix>`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2000.

Pattern: `^.+`

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)

- [AWS SDK for Ruby V3](#)

UpdateJobPriority

Service: Amazon S3 Control

Updates an existing S3 Batch Operations job's priority. For more information, see [S3 Batch Operations](#) in the *Amazon S3 User Guide*.

Permissions

To use the UpdateJobPriority operation, you must have permission to perform the `s3:UpdateJobPriority` action.

Related actions include:

- [CreateJob](#)
- [ListJobs](#)
- [DescribeJob](#)
- [UpdateJobStatus](#)

Request Syntax

```
POST /v20180820/jobs/id/priority?priority=Priority HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

id

The ID for the job whose priority you want to update.

Length Constraints: Minimum length of 5. Maximum length of 36.

Pattern: `[a-zA-Z0-9\-\_\+]`

Required: Yes

priority

The priority you want to assign to this job.

Valid Range: Minimum value of 0. Maximum value of 2147483647.

Required: Yes

[x-amz-account-id](#)

The AWS account ID associated with the S3 Batch Operations job.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<UpdateJobPriorityResult>
  <JobId>string</JobId>
  <Priority>integer</Priority>
</UpdateJobPriorityResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

[UpdateJobPriorityResult](#)

Root level tag for the UpdateJobPriorityResult parameters.

Required: Yes

[JobId](#)

The ID for the job whose priority Amazon S3 updated.

Type: String

Length Constraints: Minimum length of 5. Maximum length of 36.

Pattern: `[a-zA-Z0-9\-\_\+]`

Priority

The new priority assigned to the specified job.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 2147483647.

Errors

BadRequestException

HTTP Status Code: 400

InternalServerErrorException

HTTP Status Code: 500

NotFoundException

HTTP Status Code: 400

TooManyRequestsException

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

UpdateJobStatus

Service: Amazon S3 Control

Updates the status for the specified job. Use this operation to confirm that you want to run a job or to cancel an existing job. For more information, see [S3 Batch Operations](#) in the *Amazon S3 User Guide*.

Permissions

To use the UpdateJobStatus operation, you must have permission to perform the `s3:UpdateJobStatus` action.

Related actions include:

- [CreateJob](#)
- [ListJobs](#)
- [DescribeJob](#)
- [UpdateJobStatus](#)

Request Syntax

```
POST /v20180820/jobs/id/status?
requestedJobStatus=RequestedJobStatus&statusUpdateReason=StatusUpdateReason HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
```

URI Request Parameters

The request uses the following URI parameters.

id

The ID of the job whose status you want to update.

Length Constraints: Minimum length of 5. Maximum length of 36.

Pattern: `[a-zA-Z0-9\-\_\+]`

Required: Yes

requestedJobStatus

The status that you want to move the specified job to.

Valid Values: Cancelled | Ready

Required: Yes

statusUpdateReason

A description of the reason why you want to change the specified job's status. This field can be any string up to the maximum length.

Length Constraints: Minimum length of 1. Maximum length of 256.

x-amz-account-id

The AWS account ID associated with the S3 Batch Operations job.

Length Constraints: Maximum length of 64.

Pattern: ^\d{12}\$

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
<?xml version="1.0" encoding="UTF-8"?>
<UpdateJobStatusResult>
  <JobId>string</JobId>
  <Status>string</Status>
  <StatusUpdateReason>string</StatusUpdateReason>
</UpdateJobStatusResult>
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in XML format by the service.

UpdateJobStatusResult

Root level tag for the UpdateJobStatusResult parameters.

Required: Yes

JobId

The ID for the job whose status was updated.

Type: String

Length Constraints: Minimum length of 5. Maximum length of 36.

Pattern: [a-zA-Z0-9\-__]+

Status

The current status for the specified job.

Type: String

Valid Values: Active | Cancelled | Cancelling | Complete | Completing
| Failed | Failing | New | Paused | Pausing | Preparing | Ready |
Suspended

StatusUpdateReason

The reason that the specified job's status was updated.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Errors

BadRequestException

HTTP Status Code: 400

InternalServiceException

HTTP Status Code: 500

JobStatusException

HTTP Status Code: 400

NotFoundException

HTTP Status Code: 400

TooManyRequestsException

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

UpdateStorageLensGroup

Service: Amazon S3 Control

Updates the existing Storage Lens group.

To use this operation, you must have the permission to perform the `s3:UpdateStorageLensGroup` action. For more information about the required Storage Lens Groups permissions, see [Setting account permissions to use S3 Storage Lens groups](#).

For information about Storage Lens groups errors, see [List of Amazon S3 Storage Lens error codes](#).

Request Syntax

```
PUT /v20180820/storagelensgroup/name HTTP/1.1
Host: s3-control.amazonaws.com
x-amz-account-id: AccountId
<?xml version="1.0" encoding="UTF-8"?>
<UpdateStorageLensGroupRequest xmlns="http://awss3control.amazonaws.com/doc/2018-08-20/">
  <StorageLensGroup>
    <Filter>
      <And>
        <MatchAnyPrefix>
          <Prefix>string</Prefix>
        </MatchAnyPrefix>
        <MatchAnySuffix>
          <Suffix>string</Suffix>
        </MatchAnySuffix>
        <MatchAnyTag>
          <Tag>
            <Key>string</Key>
            <Value>string</Value>
          </Tag>
        </MatchAnyTag>
        <MatchObjectAge>
          <DaysGreaterThan>integer</DaysGreaterThan>
          <DaysLessThan>integer</DaysLessThan>
        </MatchObjectAge>
        <MatchObjectSize>
          <BytesGreaterThan>long</BytesGreaterThan>
          <BytesLessThan>long</BytesLessThan>
        </MatchObjectSize>
      </And>
    </Filter>
  </StorageLensGroup>
</UpdateStorageLensGroupRequest>
```

```

<MatchAnyPrefix>
  <Prefix>string</Prefix>
</MatchAnyPrefix>
<MatchAnySuffix>
  <Suffix>string</Suffix>
</MatchAnySuffix>
<MatchAnyTag>
  <Tag>
    <Key>string</Key>
    <Value>string</Value>
  </Tag>
</MatchAnyTag>
<MatchObjectAge>
  <DaysGreaterThan>integer</DaysGreaterThan>
  <DaysLessThan>integer</DaysLessThan>
</MatchObjectAge>
<MatchObjectSize>
  <BytesGreaterThan>long</BytesGreaterThan>
  <BytesLessThan>long</BytesLessThan>
</MatchObjectSize>
<Or>
  <MatchAnyPrefix>
    <Prefix>string</Prefix>
  </MatchAnyPrefix>
  <MatchAnySuffix>
    <Suffix>string</Suffix>
  </MatchAnySuffix>
  <MatchAnyTag>
    <Tag>
      <Key>string</Key>
      <Value>string</Value>
    </Tag>
  </MatchAnyTag>
  <MatchObjectAge>
    <DaysGreaterThan>integer</DaysGreaterThan>
    <DaysLessThan>integer</DaysLessThan>
  </MatchObjectAge>
  <MatchObjectSize>
    <BytesGreaterThan>long</BytesGreaterThan>
    <BytesLessThan>long</BytesLessThan>
  </MatchObjectSize>
</Or>
</Filter>
<Name>string</Name>

```



```
<StorageLensGroupArn>string</StorageLensGroupArn>
</StorageLensGroup>
</UpdateStorageLensGroupRequest>
```

URI Request Parameters

The request uses the following URI parameters.

[name](#)

The name of the Storage Lens group that you want to update.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\_\+]`

Required: Yes

[x-amz-account-id](#)

The AWS account ID of the Storage Lens group owner.

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Request Body

The request accepts the following data in XML format.

[UpdateStorageLensGroupRequest](#)

Root level tag for the UpdateStorageLensGroupRequest parameters.

Required: Yes

[StorageLensGroup](#)

The JSON file that contains the Storage Lens group configuration.

Type: [StorageLensGroup](#) data type

Required: Yes

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

Amazon S3 Files

The following actions are supported by Amazon S3 Files:

- [CreateAccessPoint](#)
- [CreateFileSystem](#)
- [CreateMountTarget](#)
- [DeleteAccessPoint](#)
- [DeleteFileSystem](#)

- [DeleteFileSystemPolicy](#)
- [DeleteMountTarget](#)
- [GetAccessPoint](#)
- [GetFileSystem](#)
- [GetFileSystemPolicy](#)
- [GetMountTarget](#)
- [GetSynchronizationConfiguration](#)
- [ListAccessPoints](#)
- [ListFileSystems](#)
- [ListMountTargets](#)
- [ListTagsForResource](#)
- [PutFileSystemPolicy](#)
- [PutSynchronizationConfiguration](#)
- [TagResource](#)
- [UntagResource](#)
- [UpdateMountTarget](#)

CreateAccessPoint

Service: Amazon S3 Files

Creates an S3 File System Access Point for application-specific access with POSIX user identity and root directory enforcement. Access points provide a way to manage access to shared datasets in multi-tenant scenarios.

Request Syntax

```
PUT /access-points HTTP/1.1
Content-type: application/json

{
  "clientToken": "string",
  "fileSystemId": "string",
  "posixUser": {
    "gid": number,
    "secondaryGids": [ number ],
    "uid": number
  },
  "rootDirectory": {
    "creationPermissions": {
      "ownerGid": number,
      "ownerUid": number,
      "permissions": "string"
    },
    "path": "string"
  },
  "tags": [
    {
      "key": "string",
      "value": "string"
    }
  ]
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

clientToken

A unique, case-sensitive identifier to ensure that the operation completes no more than one time. If this token matches a previous request, AWS ignores the request, but does not return an error.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: (.+)

Required: No

fileSystemId

The ID or Amazon Resource Name (ARN) of the S3 File System.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 128.

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}|fs-[0-9a-f]{17,40})

Required: Yes

posixUser

The POSIX identity with uid, gid, and secondary group IDs for user enforcement when accessing the file system through this access point.

Type: [PosixUser](#) object

Required: No

rootDirectory

The root directory path for the access point, with optional creation permissions for newly created directories.

Type: [RootDirectory](#) object

Required: No

[tags](#)

An array of key-value pairs to apply to the access point for resource tagging.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 1 item. Maximum number of 50 items.

Required: No

Response Syntax

HTTP/1.1 200

Content-type: application/json

```
{
  "accessPointArn": "string",
  "accessPointId": "string",
  "clientToken": "string",
  "fileSystemId": "string",
  "name": "string",
  "ownerId": "string",
  "posixUser": {
    "gid": number,
    "secondaryGids": [ number ],
    "uid": number
  },
  "rootDirectory": {
    "creationPermissions": {
      "ownerGid": number,
      "ownerUid": number,
      "permissions": "string"
    },
    "path": "string"
  },
  "status": "string",
  "tags": [
    {
      "key": "string",
```

```
    "value": "string"  
  }  
]  
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

accessPointArn

The Amazon Resource Name (ARN) of the access point.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 256.

Pattern: `arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}/access-point/fsap-[0-9a-f]{17,40}`

accessPointId

The ID of the access point.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 256.

Pattern: `(arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}/access-point/fsap-[0-9a-f]{17,40}|fsap-[0-9a-f]{17,40})`

clientToken

The client token that was provided in the request.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `(.+)`

fileSystemId

The ID of the S3 File System.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 128.

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}|fs-[0-9a-f]{17,40})

name

The name of the access point.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 256.

Pattern: ([\p{L}\p{Z}\p{N}_.:/+\\-@]*)

ownerId

The AWS account ID of the access point owner.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 12.

Pattern: (\d{12})|(\d{4}-\d{4}-\d{4})

posixUser

The POSIX identity configured for this access point.

Type: [PosixUser](#) object

rootDirectory

The root directory configuration for this access point.

Type: [RootDirectory](#) object

status

The current status of the access point.

Type: String

Valid Values: available | creating | deleting | deleted | error | updating

tags

The tags associated with the access point.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 1 item. Maximum number of 50 items.

Errors

ConflictException

The request conflicts with the current state of the resource. This can occur when trying to create a resource that already exists or delete a resource that is in use.

errorCode

The error code associated with the exception.

resourceId

The identifier of the resource that caused the conflict.

resourceType

The type of the resource that caused the conflict.

HTTP Status Code: 409

InternalServerError

An internal server error occurred. Retry your request.

errorCode

The error code associated with the exception.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found. Verify that the resource exists and that you have permission to access it.

errorCode

The error code associated with the exception.

HTTP Status Code: 404

ServiceQuotaExceededException

The request would exceed a service quota. Review your service quotas and either delete resources or request a quota increase.

errorCode

The error code associated with the exception.

HTTP Status Code: 402

ThrottlingException

The request was throttled. Retry your request using exponential backoff.

errorCode

The error code associated with the exception.

HTTP Status Code: 429

ValidationException

The input parameters are not valid. Check the parameter values and try again.

errorCode

The error code associated with the exception.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)

- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

CreateFileSystem

Service: Amazon S3 Files

Creates an S3 File System resource scoped to a bucket or prefix within a bucket, enabling file system access to S3 data. To create a file system, you need an S3 bucket and an IAM role that grants the service permission to access the bucket.

Request Syntax

```
PUT /file-systems HTTP/1.1
Content-type: application/json

{
  "acceptBucketWarning": boolean,
  "bucket": "string",
  "clientToken": "string",
  "kmsKeyId": "string",
  "prefix": "string",
  "roleArn": "string",
  "tags": [
    {
      "key": "string",
      "value": "string"
    }
  ]
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

acceptBucketWarning

Set to true to acknowledge and accept any warnings about the bucket configuration. If not specified, the operation may fail if there are bucket configuration warnings.

Type: Boolean

Required: No

bucket

The Amazon Resource Name (ARN) of the S3 bucket that will be accessible through the file system. The bucket must exist and be in the same AWS Region as the file system.

Type: String

Pattern: (`arn:aws[a-zA-Z0-9-]*:s3:::.+`)

Required: Yes

clientToken

A unique, case-sensitive identifier that you provide to ensure idempotent creation. Up to 64 ASCII characters are allowed. If you don't specify a client token, the AWS SDK automatically generates one.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: (`.+`)

Required: No

kmsKeyId

The ARN, key ID, or alias of the AWS KMS key to use for encryption. If not specified, the service uses a service-owned key for encryption. You can specify a KMS key using the following formats: key ID, ARN, key alias, or key alias ARN. If you use `KmsKeyId`, the file system will be encrypted.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 2048.

Pattern: (`[0-9a-f]{8}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{12}|mrk-[0-9a-f]{32}|alias/[a-zA-Z0-9/_-]+|(arn:aws[-a-z]*:kms:[a-z0-9-]+:\d{12}:((key/[0-9a-f]{8}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{12})|(key/mrk-[0-9a-f]{32})|(alias/[a-zA-Z0-9/_-]+)))`)

Required: No

[prefix](#)

An optional prefix within the S3 bucket to scope the file system access. If specified, the file system provides access only to objects with keys that begin with this prefix. If not specified, the file system provides access to the entire bucket.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 1024.

Pattern: (`|.*\/`)

Required: No

[roleArn](#)

The ARN of the IAM role that grants the S3 Files service permission to read and write data between the file system and the S3 bucket. This role must have the necessary permissions to access the specified bucket and prefix.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 2048.

Pattern: `arn:(aws[a-zA-Z-]*)?:iam::\d{12}:role/?[a-zA-Z_0-9+=,.\@\-_\/]+`

Required: Yes

[tags](#)

An array of key-value pairs to apply as tags to the file system resource. Each tag is a user-defined key-value pair. You can use tags to categorize and manage your file systems. Each key must be unique for the resource.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 1 item. Maximum number of 50 items.

Required: No

Response Syntax

```
HTTP/1.1 201
```

Content-type: application/json

```
{
  "bucket": "string",
  "clientToken": "string",
  "creationTime": number,
  "fileSystemArn": "string",
  "fileSystemId": "string",
  "kmsKeyId": "string",
  "name": "string",
  "ownerId": "string",
  "prefix": "string",
  "roleArn": "string",
  "status": "string",
  "statusMessage": "string",
  "tags": [
    {
      "key": "string",
      "value": "string"
    }
  ]
}
```

Response Elements

If the action is successful, the service sends back an HTTP 201 response.

The following data is returned in JSON format by the service.

bucket

The Amazon Resource Name (ARN) of the S3 bucket associated with the file system.

Type: String

Pattern: (arn:aws[a-zA-Z0-9-]*:s3:::.+)

clientToken

The client token used for idempotency.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: (.+)

creationTime

The time when the file system was created, in seconds since 1970-01-01T00:00:00Z (Unix epoch time).

Type: Timestamp

fileSystemArn

The ARN for the S3 file system, in the format `arn:aws:s3files:region:account-id:file-system/file-system-id`.

Type: String

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40})

fileSystemId

The ID of the file system, assigned by S3 Files. This ID is used to reference the file system in subsequent API calls.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 128.

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}|fs-[0-9a-f]{17,40})

kmsKeyId

The ARN or alias of the AWS KMS key used for encryption.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 2048.

Pattern: ([0-9a-f]{8}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{12}|mrk-[0-9a-f]{32}|alias/[a-zA-Z0-9/_-]+|(arn:aws[-a-z]*:kms:[a-zA-Z0-9-]+\d{12}:((key/[0-9a-f]{8}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{12})|(key/mrk-[0-9a-f]{32})|(alias/[a-zA-Z0-9/_-]+))))

name

The name of the file system, derived from the Name tag if present.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 256.

Pattern: (`[\\p{L}\\p{Z}\\p{N}_.: /+=\\-@]*`)

ownerId

The AWS account ID of the file system owner.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 12.

Pattern: (`\\d{12}`) | (`\\d{4}-\\d{4}-\\d{4}`)

prefix

The prefix within the S3 bucket that scopes the file system access.

Type: String

roleArn

The ARN of the IAM role used for S3 access.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 2048.

Pattern: `arn:(aws[a-zA-Z-]*)?:iam::\\d{12}:role/?[a-zA-Z_0-9+=, .@\\-_/]+`

status

The lifecycle state of the file system. Valid values are: AVAILABLE (the file system is available for use), CREATING (the file system is being created), DELETING (the file system is being deleted), DELETED (the file system has been deleted), ERROR (the file system is in an error state), or UPDATING (the file system is being updated).

Type: String

Valid Values: `available` | `creating` | `deleting` | `deleted` | `error` | `updating`

statusMessage

Additional information about the file system status. This field provides more details when the status is ERROR, or during state transitions.

Type: String

tags

The tags associated with the file system.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 1 item. Maximum number of 50 items.

Errors

ConflictException

The request conflicts with the current state of the resource. This can occur when trying to create a resource that already exists or delete a resource that is in use.

errorCode

The error code associated with the exception.

resourceId

The identifier of the resource that caused the conflict.

resourceType

The type of the resource that caused the conflict.

HTTP Status Code: 409

InternalServerError

An internal server error occurred. Retry your request.

errorCode

The error code associated with the exception.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found. Verify that the resource exists and that you have permission to access it.

errorCode

The error code associated with the exception.

HTTP Status Code: 404

ServiceQuotaExceededException

The request would exceed a service quota. Review your service quotas and either delete resources or request a quota increase.

errorCode

The error code associated with the exception.

HTTP Status Code: 402

ThrottlingException

The request was throttled. Retry your request using exponential backoff.

errorCode

The error code associated with the exception.

HTTP Status Code: 429

ValidationException

The input parameters are not valid. Check the parameter values and try again.

errorCode

The error code associated with the exception.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

CreateMountTarget

Service: Amazon S3 Files

Creates a mount target resource as an endpoint for mounting the S3 File System from compute resources in a specific Availability Zone and VPC. Mount targets provide network access to the file system.

Request Syntax

```
PUT /mount-targets HTTP/1.1
Content-type: application/json

{
  "fileSystemId": "string",
  "ipAddressType": "string",
  "ipv4Address": "string",
  "ipv6Address": "string",
  "securityGroups": [ "string" ],
  "subnetId": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

fileSystemId

The ID or Amazon Resource Name (ARN) of the S3 File System to create the mount target for.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 128.

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}|fs-[0-9a-f]{17,40})

Required: Yes

ipAddressType

The IP address type for the mount target. If not specified, IPV4_ONLY is used. The IP address type must match the IP configuration of the specified subnet.

Type: String

Valid Values: IPV4_ONLY | IPV6_ONLY | DUAL_STACK

Required: No

ipv4Address

A specific IPv4 address to assign to the mount target. If not specified and the IP address type supports IPv4, an address is automatically assigned from the subnet's available IPv4 address range. The address must be within the subnet's CIDR block and not already in use.

Type: String

Length Constraints: Minimum length of 7. Maximum length of 15.

Pattern: [0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}

Required: No

ipv6Address

A specific IPv6 address to assign to the mount target. If not specified and the IP address type supports IPv6, an address is automatically assigned from the subnet's available IPv6 address range. The address must be within the subnet's IPv6 CIDR block and not already in use.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 39.

Required: No

securityGroups

An array of VPC security group IDs to associate with the mount target's network interface. These security groups control network access to the mount target. If not specified, the default security group for the subnet's VPC is used. All security groups must belong to the same VPC as the subnet.

Type: Array of strings

Array Members: Minimum number of 0 items. Maximum number of 100 items.

Length Constraints: Minimum length of 11. Maximum length of 43.

Pattern: (sg-[0-9a-f]{8,40})

Required: No

subnetId

The ID of the subnet where the mount target will be created. The subnet must be in the same AWS Region as the file system. For file systems with regional availability, you can create mount targets in any subnet within the Region. The subnet determines the Availability Zone where the mount target will be located.

Type: String

Length Constraints: Minimum length of 15. Maximum length of 47.

Pattern: subnet-[0-9a-f]{8,40}

Required: Yes

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "availabilityZoneId": "string",
  "fileSystemId": "string",
  "ipv4Address": "string",
  "ipv6Address": "string",
  "mountTargetId": "string",
  "networkInterfaceId": "string",
  "ownerId": "string",
  "securityGroups": [ "string" ],
  "status": "string",
  "statusMessage": "string",
  "subnetId": "string",
  "vpcId": "string"
```

```
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[availabilityZoneId](#)

The unique and consistent identifier of the Availability Zone where the mount target is located. For example, use1-az1 is an Availability Zone ID for the us-east-1 AWS Region, and it has the same location in every AWS account.

Type: String

[fileSystemId](#)

The ID of the S3 File System associated with the mount target.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 128.

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+\:file-system/fs-[0-9a-f]{17,40}|fs-[0-9a-f]{17,40})

[ipv4Address](#)

The IPv4 address assigned to the mount target.

Type: String

Length Constraints: Minimum length of 7. Maximum length of 15.

Pattern: [0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}

[ipv6Address](#)

The IPv6 address assigned to the mount target.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 39.

mountTargetId

The ID of the mount target, assigned by S3 Files. This ID is used to reference the mount target in subsequent API calls.

Type: String

Length Constraints: Minimum length of 22. Maximum length of 45.

Pattern: fsmt-[0-9a-f]{17,40}

networkInterfaceId

The ID of the network interface that S3 Files created when it created the mount target. This network interface is managed by the service.

Type: String

ownerId

The AWS account ID of the mount target owner.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 12.

Pattern: (\d{12})|(\d{4}-{4}-\d{4})

securityGroups

The security groups associated with the mount target's network interface.

Type: Array of strings

Array Members: Minimum number of 0 items. Maximum number of 100 items.

Length Constraints: Minimum length of 11. Maximum length of 43.

Pattern: (sg-[0-9a-f]{8,40})

status

The lifecycle state of the mount target. Valid values are: AVAILABLE (the mount target is available for use), CREATING (the mount target is being created), DELETING (the mount target is being deleted), DELETED (the mount target has been deleted), or ERROR (the mount target is in an error state), or UPDATING (the mount target is being updated).

Type: String

Valid Values: available | creating | deleting | deleted | error | updating

statusMessage

Additional information about the mount target status. This field provides more details when the status is ERROR, or during state transitions.

Type: String

subnetId

The ID of the subnet where the mount target is located.

Type: String

Length Constraints: Minimum length of 15. Maximum length of 47.

Pattern: subnet-[0-9a-f]{8,40}

vpcId

The ID of the VPC where the mount target is located.

Type: String

Errors

ConflictException

The request conflicts with the current state of the resource. This can occur when trying to create a resource that already exists or delete a resource that is in use.

errorCode

The error code associated with the exception.

resourceId

The identifier of the resource that caused the conflict.

resourceType

The type of the resource that caused the conflict.

HTTP Status Code: 409

InternalServerErrorException

An internal server error occurred. Retry your request.

errorCode

The error code associated with the exception.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found. Verify that the resource exists and that you have permission to access it.

errorCode

The error code associated with the exception.

HTTP Status Code: 404

ServiceQuotaExceededException

The request would exceed a service quota. Review your service quotas and either delete resources or request a quota increase.

errorCode

The error code associated with the exception.

HTTP Status Code: 402

ThrottlingException

The request was throttled. Retry your request using exponential backoff.

errorCode

The error code associated with the exception.

HTTP Status Code: 429

ValidationException

The input parameters are not valid. Check the parameter values and try again.

errorCode

The error code associated with the exception.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteAccessPoint

Service: Amazon S3 Files

Deletes an S3 File System Access Point. This operation is irreversible.

Request Syntax

```
DELETE /access-points/accessPointId HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

accessPointId

The ID or Amazon Resource Name (ARN) of the access point to delete.

Length Constraints: Minimum length of 0. Maximum length of 256.

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}/access-point/fsap-[0-9a-f]{17,40}|fsap-[0-9a-f]{17,40})

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

ConflictException

The request conflicts with the current state of the resource. This can occur when trying to create a resource that already exists or delete a resource that is in use.

errorCode

The error code associated with the exception.

resourceId

The identifier of the resource that caused the conflict.

resourceType

The type of the resource that caused the conflict.

HTTP Status Code: 409

InternalServerErrorException

An internal server error occurred. Retry your request.

errorCode

The error code associated with the exception.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found. Verify that the resource exists and that you have permission to access it.

errorCode

The error code associated with the exception.

HTTP Status Code: 404

ThrottlingException

The request was throttled. Retry your request using exponential backoff.

errorCode

The error code associated with the exception.

HTTP Status Code: 429

ValidationException

The input parameters are not valid. Check the parameter values and try again.

errorCode

The error code associated with the exception.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteFileSystem

Service: Amazon S3 Files

Deletes an S3 File System. You can optionally force deletion of a file system that has pending export data.

Request Syntax

```
DELETE /file-systems/fileSystemId?forceDelete=forceDelete HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

fileSystemId

The ID or Amazon Resource Name (ARN) of the S3 File System to delete.

Length Constraints: Minimum length of 0. Maximum length of 128.

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}|fs-[0-9a-f]{17,40})

Required: Yes

forceDelete

If true, allows deletion of a file system that contains data pending export to S3. If false (the default), the deletion will fail if there is data that has not yet been exported to the S3 bucket. Use this parameter with caution as it may result in data loss.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

ConflictException

The request conflicts with the current state of the resource. This can occur when trying to create a resource that already exists or delete a resource that is in use.

errorCode

The error code associated with the exception.

resourceId

The identifier of the resource that caused the conflict.

resourceType

The type of the resource that caused the conflict.

HTTP Status Code: 409

InternalServerErrorException

An internal server error occurred. Retry your request.

errorCode

The error code associated with the exception.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found. Verify that the resource exists and that you have permission to access it.

errorCode

The error code associated with the exception.

HTTP Status Code: 404

ThrottlingException

The request was throttled. Retry your request using exponential backoff.

errorCode

The error code associated with the exception.

HTTP Status Code: 429

ValidationException

The input parameters are not valid. Check the parameter values and try again.

errorCode

The error code associated with the exception.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteFileSystemPolicy

Service: Amazon S3 Files

Deletes the IAM resource policy of an S3 File System.

Request Syntax

```
DELETE /file-systems/fileSystemId/policy HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

fileSystemId

The ID or Amazon Resource Name (ARN) of the S3 File System whose resource policy to delete.

Length Constraints: Minimum length of 0. Maximum length of 128.

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}|fs-[0-9a-f]{17,40})

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

InternalServerErrorException

An internal server error occurred. Retry your request.

errorCode

The error code associated with the exception.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found. Verify that the resource exists and that you have permission to access it.

errorCode

The error code associated with the exception.

HTTP Status Code: 404

ThrottlingException

The request was throttled. Retry your request using exponential backoff.

errorCode

The error code associated with the exception.

HTTP Status Code: 429

ValidationException

The input parameters are not valid. Check the parameter values and try again.

errorCode

The error code associated with the exception.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)

- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteMountTarget

Service: Amazon S3 Files

Deletes the specified mount target. This operation is irreversible.

Request Syntax

```
DELETE /mount-targets/mountTargetId HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

mountTargetId

The ID of the mount target to delete.

Length Constraints: Minimum length of 22. Maximum length of 45.

Pattern: fsmt-[0-9a-f]{17,40}

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

ConflictException

The request conflicts with the current state of the resource. This can occur when trying to create a resource that already exists or delete a resource that is in use.

errorCode

The error code associated with the exception.

resourceId

The identifier of the resource that caused the conflict.

resourceType

The type of the resource that caused the conflict.

HTTP Status Code: 409

InternalServerErrorException

An internal server error occurred. Retry your request.

errorCode

The error code associated with the exception.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found. Verify that the resource exists and that you have permission to access it.

errorCode

The error code associated with the exception.

HTTP Status Code: 404

ThrottlingException

The request was throttled. Retry your request using exponential backoff.

errorCode

The error code associated with the exception.

HTTP Status Code: 429

ValidationException

The input parameters are not valid. Check the parameter values and try again.

errorCode

The error code associated with the exception.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetAccessPoint

Service: Amazon S3 Files

Returns resource information for an S3 File System Access Point.

Request Syntax

```
GET /access-points/accessPointId HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

accessPointId

The ID or Amazon Resource Name (ARN) of the access point to retrieve information for.

Length Constraints: Minimum length of 0. Maximum length of 256.

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}/access-point/fsap-[0-9a-f]{17,40}|fsap-[0-9a-f]{17,40})

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "accessPointArn": "string",
  "accessPointId": "string",
  "clientToken": "string",
  "fileSystemId": "string",
  "name": "string",
  "ownerId": "string",
```

```

    "posixUser": {
      "gid": number,
      "secondaryGids": [ number ],
      "uid": number
    },
    "rootDirectory": {
      "creationPermissions": {
        "ownerGid": number,
        "ownerUid": number,
        "permissions": "string"
      },
      "path": "string"
    },
    "status": "string",
    "tags": [
      {
        "key": "string",
        "value": "string"
      }
    ]
  }
}

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

accessPointArn

The ARN of the access point.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 256.

Pattern: `arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}/access-point/fsap-[0-9a-f]{17,40}`

accessPointId

The ID of the access point.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 256.

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}/access-point/fsap-[0-9a-f]{17,40}|fsap-[0-9a-f]{17,40})

clientToken

The client token used for idempotency when the access point was created.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: (.+)

fileSystemId

The ID of the S3 File System.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 128.

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}|fs-[0-9a-f]{17,40})

name

The name of the access point.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 256.

Pattern: ([\p{L}\p{Z}\p{N}_ . : / = + \ - @] *)

ownerId

The AWS account ID of the access point owner.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 12.

Pattern: (\d{12}) | (\d{4} - \d{4} - \d{4})

posixUser

The POSIX identity configured for this access point.

Type: [PosixUser](#) object

rootDirectory

The root directory configuration for this access point.

Type: [RootDirectory](#) object

status

The current status of the access point.

Type: String

Valid Values: available | creating | deleting | deleted | error | updating

tags

The tags associated with the access point.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 1 item. Maximum number of 50 items.

Errors

InternalServerErrorException

An internal server error occurred. Retry your request.

errorCode

The error code associated with the exception.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found. Verify that the resource exists and that you have permission to access it.

errorCode

The error code associated with the exception.

HTTP Status Code: 404

ThrottlingException

The request was throttled. Retry your request using exponential backoff.

errorCode

The error code associated with the exception.

HTTP Status Code: 429

ValidationException

The input parameters are not valid. Check the parameter values and try again.

errorCode

The error code associated with the exception.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetFileSystem

Service: Amazon S3 Files

Returns resource information for the specified S3 File System including status, configuration, and metadata.

Request Syntax

```
GET /file-systems/fileSystemId HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

fileSystemId

The ID or Amazon Resource Name (ARN) of the S3 File System to retrieve information for.

Length Constraints: Minimum length of 0. Maximum length of 128.

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}|fs-[0-9a-f]{17,40})

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json
```

```
{
  "bucket": "string",
  "clientToken": "string",
  "creationTime": number,
  "fileSystemArn": "string",
  "fileSystemId": "string",
  "kmsKeyId": "string",
  "name": "string",
```

```
"ownerId": "string",
"prefix": "string",
"roleArn": "string",
"status": "string",
"statusMessage": "string",
"tags": [
  {
    "key": "string",
    "value": "string"
  }
]
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[bucket](#)

The Amazon Resource Name (ARN) of the S3 bucket.

Type: String

Pattern: (arn:aws[a-zA-Z0-9-]*:s3:::.)

[clientToken](#)

The client token used for idempotency when the file system was created.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: (.+)

[creationTime](#)

The time when the file system was created.

Type: Timestamp

[fileSystemArn](#)

The Amazon Resource Name (ARN) of the file system.

Type: String

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40})

fileSystemId

The ID of the file system.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 128.

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}|fs-[0-9a-f]{17,40})

kmsKeyId

The Amazon Resource Name (ARN) of the AWS KMS key used for encryption.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 2048.

Pattern: ([0-9a-f]{8}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{12}|mrk-[0-9a-f]{32}|alias/[a-zA-Z0-9/_-]+|(arn:aws[-a-z]*:kms:[a-z0-9-]+:\d{12}:((key/[0-9a-f]{8}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{12})|(key/mrk-[0-9a-f]{32})|(alias/[a-zA-Z0-9/_-]+))))

name

The name of the file system.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 256.

Pattern: ([\p{L}\p{Z}\p{N}_.:/+@-]*)

ownerId

The AWS account ID of the file system owner.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 12.

Pattern: `(\d{12})|(\d{4}-\d{4}-\d{4})`

[prefix](#)

The prefix in the S3 bucket that the file system provides access to.

Type: String

[roleArn](#)

The Amazon Resource Name (ARN) of the IAM role used for S3 access.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 2048.

Pattern: `arn:(aws[a-zA-Z-]*)?:iam::\d{12}:role/?[a-zA-Z_0-9+=,.\@\-_/\]+`

[status](#)

The current status of the file system.

Type: String

Valid Values: `available` | `creating` | `deleting` | `deleted` | `error` | `updating`

[statusMessage](#)

Additional information about the file system status.

Type: String

[tags](#)

The tags associated with the file system.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 1 item. Maximum number of 50 items.

Errors

InternalServerErrorException

An internal server error occurred. Retry your request.

errorCode

The error code associated with the exception.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found. Verify that the resource exists and that you have permission to access it.

errorCode

The error code associated with the exception.

HTTP Status Code: 404

ThrottlingException

The request was throttled. Retry your request using exponential backoff.

errorCode

The error code associated with the exception.

HTTP Status Code: 429

ValidationException

The input parameters are not valid. Check the parameter values and try again.

errorCode

The error code associated with the exception.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)

- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetFileSystemPolicy

Service: Amazon S3 Files

Returns the IAM resource policy of an S3 File System.

Request Syntax

```
GET /file-systems/fileSystemId/policy HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

fileSystemId

The ID or Amazon Resource Name (ARN) of the S3 File System whose resource policy to retrieve.

Length Constraints: Minimum length of 0. Maximum length of 128.

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}|fs-[0-9a-f]{17,40})

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "fileSystemId": "string",
  "policy": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

fileSystemId

The ID of the file system.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 128.

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}|fs-[0-9a-f]{17,40})

policy

The JSON-formatted resource policy for the file system.

Type: String

Errors

InternalServerError

An internal server error occurred. Retry your request.

errorCode

The error code associated with the exception.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found. Verify that the resource exists and that you have permission to access it.

errorCode

The error code associated with the exception.

HTTP Status Code: 404

ThrottlingException

The request was throttled. Retry your request using exponential backoff.

errorCode

The error code associated with the exception.

HTTP Status Code: 429

ValidationException

The input parameters are not valid. Check the parameter values and try again.

errorCode

The error code associated with the exception.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetMountTarget

Service: Amazon S3 Files

Returns detailed resource information for the specified mount target including network configuration.

Request Syntax

```
GET /mount-targets/mountTargetId HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

mountTargetId

The ID of the mount target to retrieve information for.

Length Constraints: Minimum length of 22. Maximum length of 45.

Pattern: `fsmt-[0-9a-f]{17,40}`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

```
Content-type: application/json
```

```
{
  "availabilityZoneId": "string",
  "fileSystemId": "string",
  "ipv4Address": "string",
  "ipv6Address": "string",
  "mountTargetId": "string",
  "networkInterfaceId": "string",
```



```
"ownerId": "string",  
"securityGroups": [ "string" ],  
"status": "string",  
"statusMessage": "string",  
"subnetId": "string",  
"vpcId": "string"  
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[availabilityZoneId](#)

The Availability Zone ID where the mount target is located.

Type: String

[fileSystemId](#)

The ID of the file system.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 128.

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}|fs-[0-9a-f]{17,40})

[ipv4Address](#)

The IPv4 address of the mount target.

Type: String

Length Constraints: Minimum length of 7. Maximum length of 15.

Pattern: [0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}

[ipv6Address](#)

The IPv6 address of the mount target.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 39.

mountTargetId

The ID of the mount target.

Type: String

Length Constraints: Minimum length of 22. Maximum length of 45.

Pattern: fsmt-[0-9a-f]{17,40}

networkInterfaceId

The ID of the network interface associated with the mount target.

Type: String

ownerId

The AWS account ID of the mount target owner.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 12.

Pattern: (\d{12})|(\d{4}-\d{4}-\d{4})

securityGroups

The security groups associated with the mount target.

Type: Array of strings

Array Members: Minimum number of 0 items. Maximum number of 100 items.

Length Constraints: Minimum length of 11. Maximum length of 43.

Pattern: (sg-[0-9a-f]{8,40})

status

The current status of the mount target.

Type: String

Valid Values: available | creating | deleting | deleted | error | updating

statusMessage

Additional information about the mount target status.

Type: String

subnetId

The ID of the subnet where the mount target is located.

Type: String

Length Constraints: Minimum length of 15. Maximum length of 47.

Pattern: subnet-[0-9a-f]{8,40}

vpcId

The ID of the VPC where the mount target is located.

Type: String

Errors

InternalServerError

An internal server error occurred. Retry your request.

errorCode

The error code associated with the exception.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found. Verify that the resource exists and that you have permission to access it.

errorCode

The error code associated with the exception.

HTTP Status Code: 404

ThrottlingException

The request was throttled. Retry your request using exponential backoff.

errorCode

The error code associated with the exception.

HTTP Status Code: 429

ValidationException

The input parameters are not valid. Check the parameter values and try again.

errorCode

The error code associated with the exception.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetSynchronizationConfiguration

Service: Amazon S3 Files

Returns the synchronization configuration for the specified S3 File System, including import data rules and expiration data rules.

Request Syntax

```
GET /file-systems/fileSystemId/synchronization-configuration HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

fileSystemId

The ID or Amazon Resource Name (ARN) of the S3 File System to retrieve the synchronization configuration for.

Length Constraints: Minimum length of 0. Maximum length of 128.

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}|fs-[0-9a-f]{17,40})

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "expirationDataRules": [
    {
      "daysAfterLastAccess": number
    }
  ],
```

```
"importDataRules": [  
  {  
    "prefix": "string",  
    "sizeLessThan": number,  
    "trigger": "string"  
  }  
],  
"latestVersionNumber": number  
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[expirationDataRules](#)

An array of expiration data rules that control when cached data expires from the file system.

Type: Array of [ExpirationDataRule](#) objects

Array Members: Fixed number of 1 item.

[importDataRules](#)

An array of import data rules that control how data is imported from S3 into the file system.

Type: Array of [ImportDataRule](#) objects

Array Members: Minimum number of 1 item. Maximum number of 10 items.

[latestVersionNumber](#)

The version number of the synchronization configuration. Use this value with `PutSynchronizationConfiguration` to ensure optimistic concurrency control.

Type: Integer

Errors

InternalServerErrorException

An internal server error occurred. Retry your request.

errorCode

The error code associated with the exception.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found. Verify that the resource exists and that you have permission to access it.

errorCode

The error code associated with the exception.

HTTP Status Code: 404

ThrottlingException

The request was throttled. Retry your request using exponential backoff.

errorCode

The error code associated with the exception.

HTTP Status Code: 429

ValidationException

The input parameters are not valid. Check the parameter values and try again.

errorCode

The error code associated with the exception.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)

- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListAccessPoints

Service: Amazon S3 Files

Returns resource information for all S3 File System Access Points associated with the specified S3 File System.

Request Syntax

```
GET /access-points?fileSystemId=fileSystemId&maxResults=maxResults&nextToken=nextToken  
HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

fileSystemId

The ID or Amazon Resource Name (ARN) of the S3 File System to list access points for.

Length Constraints: Minimum length of 0. Maximum length of 128.

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}|fs-[0-9a-f]{17,40})

Required: Yes

maxResults

The maximum number of access points to return in a single response.

Valid Range: Minimum value of 1. Maximum value of 1000.

nextToken

A pagination token returned from a previous call to continue listing access points.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Content-type: application/json

```
{
  "accessPoints": [
    {
      "accessPointArn": "string",
      "accessPointId": "string",
      "fileSystemId": "string",
      "name": "string",
      "ownerId": "string",
      "posixUser": {
        "gid": number,
        "secondaryGids": [ number ],
        "uid": number
      },
      "rootDirectory": {
        "creationPermissions": {
          "ownerGid": number,
          "ownerUid": number,
          "permissions": "string"
        },
        "path": "string"
      },
      "status": "string"
    }
  ],
  "nextToken": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

accessPoints

An array of access point descriptions.

Type: Array of [ListAccessPointsDescription](#) objects

nextToken

A pagination token to use in a subsequent request if more results are available.

Type: String

Errors

InternalServerErrorException

An internal server error occurred. Retry your request.

errorCode

The error code associated with the exception.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found. Verify that the resource exists and that you have permission to access it.

errorCode

The error code associated with the exception.

HTTP Status Code: 404

ThrottlingException

The request was throttled. Retry your request using exponential backoff.

errorCode

The error code associated with the exception.

HTTP Status Code: 429

ValidationException

The input parameters are not valid. Check the parameter values and try again.

errorCode

The error code associated with the exception.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListFileSystems

Service: Amazon S3 Files

Returns a list of all S3 File Systems owned by the account with optional filtering by bucket.

Request Syntax

```
GET /file-systems?bucket=bucket&maxResults=maxResults&nextToken=nextToken HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

bucket

Optional filter to list only file systems associated with the specified S3 bucket Amazon Resource Name (ARN). If provided, only file systems that provide access to this bucket will be returned in the response.

Pattern: (arn:aws[a-zA-Z0-9-]*:s3:::.+)

maxResults

The maximum number of file systems to return in a single response. If not specified, up to 100 file systems are returned.

Valid Range: Minimum value of 1. Maximum value of 100.

nextToken

A pagination token returned from a previous call to continue listing file systems.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200  
Content-type: application/json
```

```
{
  "fileSystems": [
    {
      "bucket": "string",
      "creationTime": number,
      "fileSystemArn": "string",
      "fileSystemId": "string",
      "name": "string",
      "ownerId": "string",
      "roleArn": "string",
      "status": "string",
      "statusMessage": "string"
    }
  ],
  "nextToken": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

fileSystems

An array of file system descriptions.

Type: Array of [ListFileSystemsDescription](#) objects

nextToken

A pagination token to use in a subsequent request if more results are available.

Type: String

Errors

InternalServerError

An internal server error occurred. Retry your request.

errorCode

The error code associated with the exception.

HTTP Status Code: 500

ThrottlingException

The request was throttled. Retry your request using exponential backoff.

errorCode

The error code associated with the exception.

HTTP Status Code: 429

ValidationException

The input parameters are not valid. Check the parameter values and try again.

errorCode

The error code associated with the exception.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListMountTargets

Service: Amazon S3 Files

Returns resource information for all mount targets with optional filtering by file system, access point, and VPC.

Request Syntax

```
GET /mount-targets?  
accessPointId=accessPointId&fileSystemId=fileSystemId&maxResults=maxResults&nextToken=nextToken  
HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

accessPointId

Optional filter to list only mount targets associated with the specified access point ID or Amazon Resource Name (ARN).

Length Constraints: Minimum length of 0. Maximum length of 256.

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}/access-point/fsap-[0-9a-f]{17,40}|fsap-[0-9a-f]{17,40})

fileSystemId

Optional filter to list only mount targets associated with the specified S3 File System ID or Amazon Resource Name (ARN). If provided, only mount targets for this file system will be returned in the response.

Length Constraints: Minimum length of 0. Maximum length of 128.

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}|fs-[0-9a-f]{17,40})

maxResults

The maximum number of mount targets to return in a single response.

Valid Range: Minimum value of 1. Maximum value of 100.

[nextToken](#)

A pagination token returned from a previous call to continue listing mount targets.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "mountTargets": [
    {
      "availabilityZoneId": "string",
      "fileSystemId": "string",
      "ipv4Address": "string",
      "ipv6Address": "string",
      "mountTargetId": "string",
      "networkInterfaceId": "string",
      "ownerId": "string",
      "status": "string",
      "statusMessage": "string",
      "subnetId": "string",
      "vpcId": "string"
    }
  ],
  "nextToken": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[mountTargets](#)

An array of mount target descriptions.

Type: Array of [ListMountTargetsDescription](#) objects

nextToken

A pagination token to use in a subsequent request if more results are available.

Type: String

Errors

InternalServerError

An internal server error occurred. Retry your request.

errorCode

The error code associated with the exception.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found. Verify that the resource exists and that you have permission to access it.

errorCode

The error code associated with the exception.

HTTP Status Code: 404

ThrottlingException

The request was throttled. Retry your request using exponential backoff.

errorCode

The error code associated with the exception.

HTTP Status Code: 429

ValidationException

The input parameters are not valid. Check the parameter values and try again.

errorCode

The error code associated with the exception.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListTagsForResource

Service: Amazon S3 Files

Lists all tags for S3 Files resources.

Request Syntax

```
GET /resource-tags/resourceId?MaxResults=maxResults&NextToken=nextToken HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

maxResults

The maximum number of tags to return in a single response.

Valid Range: Minimum value of 1. Maximum value of 50.

nextToken

A pagination token returned from a previous call to continue listing tags.

resourceId

The ID or Amazon Resource Name (ARN) of the resource to list tags for.

Length Constraints: Minimum length of 0. Maximum length of 256.

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}
(/access-point/fsap-[0-9a-f]{17,40})?|fs(ap)?-[0-9a-f]{17,40})

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

```
Content-type: application/json
```

```
{
  "nextToken": "string",
  "tags": [
    {
      "key": "string",
      "value": "string"
    }
  ]
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

nextToken

A pagination token to use in a subsequent request if more results are available.

Type: String

tags

An array of tags associated with the resource.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 1 item. Maximum number of 50 items.

Errors

InternalServerError

An internal server error occurred. Retry your request.

errorCode

The error code associated with the exception.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found. Verify that the resource exists and that you have permission to access it.

errorCode

The error code associated with the exception.

HTTP Status Code: 404

ThrottlingException

The request was throttled. Retry your request using exponential backoff.

errorCode

The error code associated with the exception.

HTTP Status Code: 429

ValidationException

The input parameters are not valid. Check the parameter values and try again.

errorCode

The error code associated with the exception.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)

- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutFileSystemPolicy

Service: Amazon S3 Files

Creates or replaces the IAM resource policy for an S3 File System to control access permissions.

Request Syntax

```
PUT /file-systems/fileSystemId/policy HTTP/1.1
Content-type: application/json

{
  "policy": "string"
}
```

URI Request Parameters

The request uses the following URI parameters.

fileSystemId

The ID or Amazon Resource Name (ARN) of the S3 File System to apply the resource policy to.

Length Constraints: Minimum length of 0. Maximum length of 128.

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}|fs-[0-9a-f]{17,40})

Required: Yes

Request Body

The request accepts the following data in JSON format.

policy

The JSON-formatted resource policy to apply to the file system. The policy defines the permissions for accessing the file system. The policy must be a valid JSON document that follows IAM policy syntax.

Type: String

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

InternalServerError

An internal server error occurred. Retry your request.

errorCode

The error code associated with the exception.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found. Verify that the resource exists and that you have permission to access it.

errorCode

The error code associated with the exception.

HTTP Status Code: 404

ThrottlingException

The request was throttled. Retry your request using exponential backoff.

errorCode

The error code associated with the exception.

HTTP Status Code: 429

ValidationException

The input parameters are not valid. Check the parameter values and try again.

errorCode

The error code associated with the exception.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutSynchronizationConfiguration

Service: Amazon S3 Files

Creates or updates the synchronization configuration for the specified S3 File System, including import data rules and expiration data rules.

Request Syntax

```
PUT /file-systems/fileSystemId/synchronization-configuration HTTP/1.1
```

Content-type: application/json

```
{
  "expirationDataRules": [
    {
      "daysAfterLastAccess": number
    }
  ],
  "importDataRules": [
    {
      "prefix": "string",
      "sizeLessThan": number,
      "trigger": "string"
    }
  ],
  "latestVersionNumber": number
}
```

URI Request Parameters

The request uses the following URI parameters.

fileSystemId

The ID or Amazon Resource Name (ARN) of the S3 File System to configure synchronization for.

Length Constraints: Minimum length of 0. Maximum length of 128.

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}|fs-[0-9a-f]{17,40})

Required: Yes

Request Body

The request accepts the following data in JSON format.

expirationDataRules

An array of expiration data rules that control when cached data expires from the file system.

Type: Array of [ExpirationDataRule](#) objects

Array Members: Fixed number of 1 item.

Required: Yes

importDataRules

An array of import data rules that control how data is imported from S3 into the file system.

Type: Array of [ImportDataRule](#) objects

Array Members: Minimum number of 1 item. Maximum number of 10 items.

Required: Yes

latestVersionNumber

The version number of the current synchronization configuration. Omit this value when creating a synchronization configuration for the first time. For subsequent updates, provide this value for optimistic concurrency control. If the version number does not match the current configuration, the request fails with a `ConflictException`.

Type: Integer

Required: No

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

ConflictException

The request conflicts with the current state of the resource. This can occur when trying to create a resource that already exists or delete a resource that is in use.

errorCode

The error code associated with the exception.

resourceId

The identifier of the resource that caused the conflict.

resourceType

The type of the resource that caused the conflict.

HTTP Status Code: 409

InternalServerErrorException

An internal server error occurred. Retry your request.

errorCode

The error code associated with the exception.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found. Verify that the resource exists and that you have permission to access it.

errorCode

The error code associated with the exception.

HTTP Status Code: 404

ThrottlingException

The request was throttled. Retry your request using exponential backoff.

errorCode

The error code associated with the exception.

HTTP Status Code: 429

ValidationException

The input parameters are not valid. Check the parameter values and try again.

errorCode

The error code associated with the exception.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

TagResource

Service: Amazon S3 Files

Creates tags for S3 Files resources using standard AWS tagging APIs.

Request Syntax

```
POST /resource-tags/resourceId HTTP/1.1
Content-type: application/json
```

```
{
  "tags": [
    {
      "key": "string",
      "value": "string"
    }
  ]
}
```

URI Request Parameters

The request uses the following URI parameters.

resourceId

The ID or Amazon Resource Name (ARN) of the resource to add tags to.

Length Constraints: Minimum length of 0. Maximum length of 256.

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}(/access-point/fsap-[0-9a-f]{17,40})?|fs(ap)?-[0-9a-f]{17,40})

Required: Yes

Request Body

The request accepts the following data in JSON format.

tags

An array of key-value pairs to add as tags to the resource.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 1 item. Maximum number of 50 items.

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

InternalServerError

An internal server error occurred. Retry your request.

errorCode

The error code associated with the exception.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found. Verify that the resource exists and that you have permission to access it.

errorCode

The error code associated with the exception.

HTTP Status Code: 404

ThrottlingException

The request was throttled. Retry your request using exponential backoff.

errorCode

The error code associated with the exception.

HTTP Status Code: 429

ValidationException

The input parameters are not valid. Check the parameter values and try again.

errorCode

The error code associated with the exception.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

UntagResource

Service: Amazon S3 Files

Removes tags from S3 Files resources.

Request Syntax

```
DELETE /resource-tags/resourceId?tagKeys=tagKeys HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

resourceId

The ID or Amazon Resource Name (ARN) of the resource to remove tags from.

Length Constraints: Minimum length of 0. Maximum length of 256.

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}(/access-point/fsap-[0-9a-f]{17,40})?|fs(ap)?-[0-9a-f]{17,40})

Required: Yes

tagKeys

An array of tag keys to remove from the resource.

Array Members: Minimum number of 1 item. Maximum number of 50 items.

Length Constraints: Minimum length of 1. Maximum length of 128.

Pattern: ([\p{L}\p{Z}\p{N}_.:/+\\-@]+)

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

InternalServerError

An internal server error occurred. Retry your request.

errorCode

The error code associated with the exception.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found. Verify that the resource exists and that you have permission to access it.

errorCode

The error code associated with the exception.

HTTP Status Code: 404

ThrottlingException

The request was throttled. Retry your request using exponential backoff.

errorCode

The error code associated with the exception.

HTTP Status Code: 429

ValidationException

The input parameters are not valid. Check the parameter values and try again.

errorCode

The error code associated with the exception.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

UpdateMountTarget

Service: Amazon S3 Files

Updates the mount target resource, specifically security group configurations.

Request Syntax

```
PUT /mount-targets/mountTargetId HTTP/1.1
Content-type: application/json

{
  "securityGroups": [ "string" ]
}
```

URI Request Parameters

The request uses the following URI parameters.

mountTargetId

The ID of the mount target to update.

Length Constraints: Minimum length of 22. Maximum length of 45.

Pattern: fsmt-[0-9a-f]{17,40}

Required: Yes

Request Body

The request accepts the following data in JSON format.

securityGroups

An array of VPC security group IDs to associate with the mount target's network interface. This replaces the existing security groups. All security groups must belong to the same VPC as the mount target's subnet.

Type: Array of strings

Array Members: Minimum number of 0 items. Maximum number of 100 items.

Length Constraints: Minimum length of 11. Maximum length of 43.

Pattern: (sg-[0-9a-f]{8,40})

Required: Yes

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "availabilityZoneId": "string",
  "fileSystemId": "string",
  "ipv4Address": "string",
  "ipv6Address": "string",
  "mountTargetId": "string",
  "networkInterfaceId": "string",
  "ownerId": "string",
  "securityGroups": [ "string" ],
  "status": "string",
  "statusMessage": "string",
  "subnetId": "string",
  "vpcId": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

availabilityZoneId

The Availability Zone ID where the mount target is located.

Type: String

fileSystemId

The ID of the S3 File System.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 128.

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}|fs-[0-9a-f]{17,40})

ipv4Address

The IPv4 address of the mount target.

Type: String

Length Constraints: Minimum length of 7. Maximum length of 15.

Pattern: [0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}

ipv6Address

The IPv6 address of the mount target.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 39.

mountTargetId

The ID of the mount target.

Type: String

Length Constraints: Minimum length of 22. Maximum length of 45.

Pattern: fsmt-[0-9a-f]{17,40}

networkInterfaceId

The ID of the network interface associated with the mount target.

Type: String

ownerId

The AWS account ID of the mount target owner.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 12.

Pattern: (\d{12})|(\d{4}-\d{4}-\d{4})

securityGroups

The security groups associated with the mount target.

Type: Array of strings

Array Members: Minimum number of 0 items. Maximum number of 100 items.

Length Constraints: Minimum length of 11. Maximum length of 43.

Pattern: (sg-[0-9a-f]{8,40})

status

The current status of the mount target.

Type: String

Valid Values: available | creating | deleting | deleted | error | updating

statusMessage

Additional information about the mount target status.

Type: String

subnetId

The ID of the subnet where the mount target is located.

Type: String

Length Constraints: Minimum length of 15. Maximum length of 47.

Pattern: subnet-[0-9a-f]{8,40}

vpcId

The ID of the VPC where the mount target is located.

Type: String

Errors

InternalServerError

An internal server error occurred. Retry your request.

errorCode

The error code associated with the exception.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found. Verify that the resource exists and that you have permission to access it.

errorCode

The error code associated with the exception.

HTTP Status Code: 404

ThrottlingException

The request was throttled. Retry your request using exponential backoff.

errorCode

The error code associated with the exception.

HTTP Status Code: 429

ValidationException

The input parameters are not valid. Check the parameter values and try again.

errorCode

The error code associated with the exception.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)

- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

Amazon S3 on Outposts

The following actions are supported by Amazon S3 on Outposts:

- [CreateEndpoint](#)
- [DeleteEndpoint](#)
- [ListEndpoints](#)
- [ListOutpostsWithS3](#)
- [ListSharedEndpoints](#)

CreateEndpoint

Service: Amazon S3 on Outposts

Creates an endpoint and associates it with the specified Outpost.

Note

It can take up to 5 minutes for this action to finish.

Related actions include:

- [DeleteEndpoint](#)
- [ListEndpoints](#)

Request Syntax

```
POST /S3Outposts/CreateEndpoint HTTP/1.1
```

```
Content-type: application/json
```

```
{
  "AccessType": "string",
  "CustomerOwnedIpv4Pool": "string",
  "OutpostId": "string",
  "SecurityGroupId": "string",
  "SubnetId": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

[AccessType](#)

The type of access for the network connectivity for the Amazon S3 on Outposts endpoint. To use the AWS VPC, choose Private. To use the endpoint with an on-premises network, choose

CustomerOwnedIp. If you choose CustomerOwnedIp, you must also provide the customer-owned IP address pool (CoIP pool).

 Note

Private is the default access type value.

Type: String

Valid Values: Private | CustomerOwnedIp

Required: No

CustomerOwnedIpv4Pool

The ID of the customer-owned IPv4 address pool (CoIP pool) for the endpoint. IP addresses are allocated from this pool for the endpoint.

Type: String

Pattern: ^ipv4pool-coip-([0-9a-f]{17})\$

Required: No

OutpostId

The ID of the AWS Outposts.

Type: String

Pattern: ^(op-[a-f0-9]{17}|\d{12}|ec2)\$

Required: Yes

SecurityGroupId

The ID of the security group to use with the endpoint.

Type: String

Pattern: ^sg-([0-9a-f]{8}|[0-9a-f]{17})\$

Required: Yes

SubnetId

The ID of the subnet in the selected VPC. The endpoint subnet must belong to the Outpost that has Amazon S3 on Outposts provisioned.

Type: String

Pattern: ^subnet-([0-9a-f]{8}|[0-9a-f]{17})\$

Required: Yes

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "EndpointArn": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

EndpointArn

The Amazon Resource Name (ARN) of the endpoint.

Type: String

Pattern: ^arn:(aws|aws-cn|aws-us-gov|aws-iso|aws-iso-b):s3-outposts:[a-z\-\0-9]*:[0-9]{12}:outpost/(op-[a-f0-9]{17}|ec2)/endpoint/[a-zA-Z0-9]{19}\$

Errors

AccessDeniedException

Access was denied for this action.

HTTP Status Code: 403

ConflictException

There was a conflict with this action, and it could not be completed.

HTTP Status Code: 409

InternalServerErrorException

There was an exception with the internal server.

HTTP Status Code: 500

OutpostOfflineException

The service link connection to your Outposts home Region is down. Check your connection and try again.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource was not found.

HTTP Status Code: 404

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 429

ValidationException

There was an exception validating this data.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)

- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteEndpoint

Service: Amazon S3 on Outposts

Deletes an endpoint.

Note

It can take up to 5 minutes for this action to finish.

Related actions include:

- [CreateEndpoint](#)
- [ListEndpoints](#)

Request Syntax

```
DELETE /S3Outposts/DeleteEndpoint?endpointId=EndpointId&outpostId=OutpostId HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

EndpointId

The ID of the endpoint.

Pattern: `^[a-zA-Z0-9]{19}$`

Required: Yes

OutpostId

The ID of the AWS Outposts.

Pattern: `^(op-[a-f0-9]{17}|\d{12}|ec2)$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

AccessDeniedException

Access was denied for this action.

HTTP Status Code: 403

InternalServerErrorException

There was an exception with the internal server.

HTTP Status Code: 500

OutpostOfflineException

The service link connection to your Outposts home Region is down. Check your connection and try again.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource was not found.

HTTP Status Code: 404

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 429

ValidationException

There was an exception validating this data.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListEndpoints

Service: Amazon S3 on Outposts

Lists endpoints associated with the specified Outpost.

Related actions include:

- [CreateEndpoint](#)
- [DeleteEndpoint](#)

Request Syntax

```
GET /S3Outposts/ListEndpoints?maxResults=MaxResults&nextToken=NextToken HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

[MaxResults](#)

The maximum number of endpoints that will be returned in the response.

Valid Range: Minimum value of 0. Maximum value of 100.

[NextToken](#)

If a previous response from this operation included a NextToken value, provide that value here to retrieve the next page of results.

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `^[A-Za-z0-9\+\:\.\=\?\#\_\-]+$`

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json
```

```

{
  "Endpoints": [
    {
      "AccessType": "string",
      "CidrBlock": "string",
      "CreationTime": number,
      "CustomerOwnedIpv4Pool": "string",
      "EndpointArn": "string",
      "FailedReason": {
        "ErrorCode": "string",
        "Message": "string"
      },
      "NetworkInterfaces": [
        {
          "NetworkInterfaceId": "string"
        }
      ],
      "OutpostsId": "string",
      "SecurityGroupId": "string",
      "Status": "string",
      "SubnetId": "string",
      "VpcId": "string"
    }
  ],
  "NextToken": "string"
}

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[Endpoints](#)

The list of endpoints associated with the specified Outpost.

Type: Array of [Endpoint](#) objects

[NextToken](#)

If the number of endpoints associated with the specified Outpost exceeds `MaxResults`, you can include this value in subsequent calls to this operation to retrieve more results.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `^[A-Za-z0-9\+\:\.\=\?\#\_\-]+$`

Errors

AccessDeniedException

Access was denied for this action.

HTTP Status Code: 403

InternalServerErrorException

There was an exception with the internal server.

HTTP Status Code: 500

ResourceNotFoundException

The requested resource was not found.

HTTP Status Code: 404

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 429

ValidationException

There was an exception validating this data.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)

- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListOutpostsWithS3

Service: Amazon S3 on Outposts

Lists the Outposts with S3 on Outposts capacity for your AWS account. Includes S3 on Outposts that you have access to as the Outposts owner, or as a shared user from Resource Access Manager (RAM).

Request Syntax

```
GET /S3Outposts/ListOutpostsWithS3?maxResults=MaxResults&nextToken=NextToken HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

MaxResults

The maximum number of Outposts to return. The limit is 100.

Valid Range: Minimum value of 0. Maximum value of 100.

NextToken

When you can get additional results from the ListOutpostsWithS3 call, a NextToken parameter is returned in the output. You can then pass in a subsequent command to the NextToken parameter to continue listing additional Outposts.

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `^[A-Za-z0-9\+\:\\/\=\?\#\_\-]+$`

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "NextToken": "string",
```

```
"Outposts": [  
  {  
    "CapacityInBytes": number,  
    "OutpostArn": "string",  
    "OutpostId": "string",  
    "OwnerId": "string",  
    "S3OutpostArn": "string"  
  }  
]
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[NextToken](#)

Returns a token that you can use to call `ListOutpostsWithS3` again and receive additional results, if there are any.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `^[A-Za-z0-9\+\:\.\\/\=\?\#\_\-]+$`

[Outposts](#)

Returns the list of Outposts that have the following characteristics:

- outposts that have S3 provisioned
- outposts that are Active (not pending any provisioning nor decommissioned)
- outposts to which the the calling AWS account has access

Type: Array of [Outpost](#) objects

Errors

AccessDeniedException

Access was denied for this action.

HTTP Status Code: 403

InternalServerErrorException

There was an exception with the internal server.

HTTP Status Code: 500

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 429

ValidationException

There was an exception validating this data.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListSharedEndpoints

Service: Amazon S3 on Outposts

Lists all endpoints associated with an Outpost that has been shared by AWS Resource Access Manager (RAM).

Related actions include:

- [CreateEndpoint](#)
- [DeleteEndpoint](#)

Request Syntax

```
GET /S3Outposts/ListSharedEndpoints?  
maxResults=MaxResults&nextToken=NextToken&outpostId=OutpostId HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

[MaxResults](#)

The maximum number of endpoints that will be returned in the response.

Valid Range: Minimum value of 0. Maximum value of 100.

[NextToken](#)

If a previous response from this operation included a NextToken value, you can provide that value here to retrieve the next page of results.

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `^[A-Za-z0-9\+\:\.\=\?\#\_\-]+$`

[OutpostId](#)

The ID of the AWS Outpost.

Pattern: `^(op-[a-f0-9]{17}|\d{12}|ec2)$`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "Endpoints": [
    {
      "AccessType": "string",
      "CidrBlock": "string",
      "CreationTime": number,
      "CustomerOwnedIpv4Pool": "string",
      "EndpointArn": "string",
      "FailedReason": {
        "ErrorCode": "string",
        "Message": "string"
      },
      "NetworkInterfaces": [
        {
          "NetworkInterfaceId": "string"
        }
      ],
      "OutpostsId": "string",
      "SecurityGroupId": "string",
      "Status": "string",
      "SubnetId": "string",
      "VpcId": "string"
    }
  ],
  "NextToken": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

Endpoints

The list of endpoints associated with the specified Outpost that have been shared by AWS Resource Access Manager (RAM).

Type: Array of [Endpoint](#) objects

NextToken

If the number of endpoints associated with the specified Outpost exceeds `MaxResults`, you can include this value in subsequent calls to this operation to retrieve more results.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `^[A-Za-z0-9\+\:\.\=\?\#\_\-]+\`

Errors

AccessDeniedException

Access was denied for this action.

HTTP Status Code: 403

InternalServerErrorException

There was an exception with the internal server.

HTTP Status Code: 500

ResourceNotFoundException

The requested resource was not found.

HTTP Status Code: 404

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 429

ValidationException

There was an exception validating this data.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

Amazon S3 Tables

The following actions are supported by Amazon S3 Tables:

- [CreateNamespace](#)
- [CreateTable](#)
- [CreateTableBucket](#)
- [DeleteNamespace](#)
- [DeleteTable](#)
- [DeleteTableBucket](#)
- [DeleteTableBucketEncryption](#)
- [DeleteTableBucketMetricsConfiguration](#)

- [DeleteTableBucketPolicy](#)
- [DeleteTableBucketReplication](#)
- [DeleteTablePolicy](#)
- [DeleteTableReplication](#)
- [GetNamespace](#)
- [GetTable](#)
- [GetTableBucket](#)
- [GetTableBucketEncryption](#)
- [GetTableBucketMaintenanceConfiguration](#)
- [GetTableBucketMetricsConfiguration](#)
- [GetTableBucketPolicy](#)
- [GetTableBucketReplication](#)
- [GetTableBucketStorageClass](#)
- [GetTableEncryption](#)
- [GetTableMaintenanceConfiguration](#)
- [GetTableMaintenanceJobStatus](#)
- [GetTableMetadataLocation](#)
- [GetTablePolicy](#)
- [GetTableRecordExpirationConfiguration](#)
- [GetTableRecordExpirationJobStatus](#)
- [GetTableReplication](#)
- [GetTableReplicationStatus](#)
- [GetTableStorageClass](#)
- [ListNamespaces](#)
- [ListTableBuckets](#)
- [ListTables](#)
- [ListTagsForResource](#)
- [PutTableBucketEncryption](#)
- [PutTableBucketMaintenanceConfiguration](#)
- [PutTableBucketMetricsConfiguration](#)

- [PutTableBucketPolicy](#)
- [PutTableBucketReplication](#)
- [PutTableBucketStorageClass](#)
- [PutTableMaintenanceConfiguration](#)
- [PutTablePolicy](#)
- [PutTableRecordExpirationConfiguration](#)
- [PutTableReplication](#)
- [RenameTable](#)
- [TagResource](#)
- [UntagResource](#)
- [UpdateTableMetadataLocation](#)

CreateNamespace

Service: Amazon S3 Tables

Creates a namespace. A namespace is a logical grouping of tables within your table bucket, which you can use to organize tables. For more information, see [Create a namespace](#) in the *Amazon Simple Storage Service User Guide*.

Permissions

You must have the `s3tables:CreateNamespace` permission to use this operation.

Request Syntax

```
PUT /namespaces/tableBucketARN HTTP/1.1
Content-type: application/json

{
  "namespace": [ "string" ]
}
```

URI Request Parameters

The request uses the following URI parameters.

[tableBucketARN](#)

The Amazon Resource Name (ARN) of the table bucket to create the namespace in.

Pattern: (`arn:aws``[-a-z0-9]*``:[a-z0-9]+``:[-a-z0-9]*``:[0-9]{12}``:bucket/``[a-z0-9_-]{3,63}`)

Required: Yes

Request Body

The request accepts the following data in JSON format.

[namespace](#)

A name for the namespace.

Type: Array of strings

Array Members: Fixed number of 1 item.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "namespace": [ "string" ],
  "tableBucketARN": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

namespace

The name of the namespace.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket the namespace was created in.

Type: String

Pattern: `(arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})`

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)

- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

CreateTable

Service: Amazon S3 Tables

Creates a new table associated with the given namespace in a table bucket. For more information, see [Creating an Amazon S3 table](#) in the *Amazon Simple Storage Service User Guide*.

Permissions

- You must have the `s3tables:CreateTable` permission to use this operation.
- If you use this operation with the optional `metadata` request parameter you must have the `s3tables:PutTableData` permission.
- If you use this operation with the optional `encryptionConfiguration` request parameter you must have the `s3tables:PutTableEncryption` permission.
- If you use this operation with the `storageClassConfiguration` request parameter, you must have the `s3tables:PutTableStorageClass` permission.
- To create a table with tags, you must have the `s3tables:TagResource` permission in addition to `s3tables:CreateTable` permission.

Note

Additionally, If you choose SSE-KMS encryption you must grant the S3 Tables maintenance principal access to your KMS key. For more information, see [Permissions requirements for S3 Tables SSE-KMS encryption](#).

Request Syntax

```
PUT /tables/tableBucketARN/namespace HTTP/1.1
```

```
Content-type: application/json
```

```
{
  "encryptionConfiguration": {
    "kmsKeyArn": "string",
    "sseAlgorithm": "string"
  },
  "format": "string",
  "metadata": { ... },
  "name": "string",
  "storageClassConfiguration": {
```

```
    "storageClass": "string"
  },
  "tags": {
    "string" : "string"
  }
}
```

URI Request Parameters

The request uses the following URI parameters.

namespace

The namespace to associated with the table.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket to create the table in.

Pattern: `(arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})`

Required: Yes

Request Body

The request accepts the following data in JSON format.

encryptionConfiguration

The encryption configuration to use for the table. This configuration specifies the encryption algorithm and, if using SSE-KMS, the KMS key to use for encrypting the table.

Note

If you choose SSE-KMS encryption you must grant the S3 Tables maintenance principal access to your KMS key. For more information, see [Permissions requirements for S3 Tables SSE-KMS encryption](#).

Type: [EncryptionConfiguration](#) object

Required: No

format

The format for the table.

Type: String

Valid Values: ICEBERG

Required: Yes

metadata

The metadata for the table.

Type: [TableMetadata](#) object

Note: This object is a Union. Only one member of this object can be specified or returned.

Required: No

name

The name for the table.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

storageClassConfiguration

The storage class configuration for the table. If not specified, the table inherits the storage class configuration from its table bucket. Specify this parameter to override the bucket's default storage class for this table.

Type: [StorageClassConfiguration](#) object

Required: No

tags

A map of user-defined tags that you would like to apply to the table that you are creating. A tag is a key-value pair that you apply to your resources. Tags can help you organize, track costs for, and control access to resources. For more information, see [Tagging for cost allocation or attribute-based access control \(ABAC\)](#).

Note

You must have the `s3tables:TagResource` permission in addition to `s3tables:CreateTable` permission to create a table with tags.

Type: String to string map

Key Length Constraints: Minimum length of 1. Maximum length of 128.

Value Length Constraints: Minimum length of 0. Maximum length of 256.

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "tableARN": "string",
  "versionToken": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

tableARN

The Amazon Resource Name (ARN) of the table.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `(arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63}/table/[a-zA-Z0-9-_{1,255})`

versionToken

The version token of the table.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

CreateTableBucket

Service: Amazon S3 Tables

Creates a table bucket. For more information, see [Creating a table bucket](#) in the *Amazon Simple Storage Service User Guide*.

Permissions

- You must have the `s3tables:CreateTableBucket` permission to use this operation.
- If you use this operation with the optional `encryptionConfiguration` parameter you must have the `s3tables:PutTableBucketEncryption` permission.
- If you use this operation with the `storageClassConfiguration` request parameter, you must have the `s3tables:PutTableBucketStorageClass` permission.
- To create a table bucket with tags, you must have the `s3tables:TagResource` permission in addition to `s3tables:CreateTableBucket` permission.

Request Syntax

```
PUT /buckets HTTP/1.1
Content-type: application/json

{
  "encryptionConfiguration": {
    "kmsKeyArn": "string",
    "sseAlgorithm": "string"
  },
  "name": "string",
  "storageClassConfiguration": {
    "storageClass": "string"
  },
  "tags": {
    "string" : "string"
  }
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

encryptionConfiguration

The encryption configuration to use for the table bucket. This configuration specifies the default encryption settings that will be applied to all tables created in this bucket unless overridden at the table level. The configuration includes the encryption algorithm and, if using SSE-KMS, the KMS key to use.

Type: [EncryptionConfiguration](#) object

Required: No

name

The name for the table bucket.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Pattern: [0-9a-z-]*

Required: Yes

storageClassConfiguration

The default storage class configuration for the table bucket. This configuration will be applied to all new tables created in this bucket unless overridden at the table level. If not specified, the service default storage class will be used.

Type: [StorageClassConfiguration](#) object

Required: No

tags

A map of user-defined tags that you would like to apply to the table bucket that you are creating. A tag is a key-value pair that you apply to your resources. Tags can help you organize and control access to resources. For more information, see [Tagging for cost allocation or attribute-based access control \(ABAC\)](#).

Note

You must have the `s3tables:TagResource` permission in addition to `s3tables:CreateTableBucket` permission to create a table bucket with tags.

Type: String to string map

Key Length Constraints: Minimum length of 1. Maximum length of 128.

Value Length Constraints: Minimum length of 0. Maximum length of 256.

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json
```

```
{
  "arn": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[arn](#)

The Amazon Resource Name (ARN) of the table bucket.

Type: String

Pattern: (`arn:aws``[-a-z0-9]*``:[a-z0-9]+``:[-a-z0-9]*``:[0-9]{12}``:bucket/``[a-z0-9_-]{3,63}`)

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)

- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteNamespace

Service: Amazon S3 Tables

Deletes a namespace. For more information, see [Delete a namespace](#) in the *Amazon Simple Storage Service User Guide*.

Permissions

You must have the `s3tables:DeleteNamespace` permission to use this operation.

Request Syntax

```
DELETE /namespaces/tableBucketARN/namespace HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

[namespace](#)

The name of the namespace.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

[tableBucketARN](#)

The Amazon Resource Name (ARN) of the table bucket associated with the namespace.

Pattern: `(arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteTable

Service: Amazon S3 Tables

Deletes a table. For more information, see [Deleting an Amazon S3 table](#) in the *Amazon Simple Storage Service User Guide*.

Permissions

You must have the `s3tables:DeleteTable` permission to use this operation.

Request Syntax

```
DELETE /tables/tableBucketARN/namespace/name?versionToken=versionToken HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the table.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

namespace

The namespace associated with the table.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket that contains the table.

Pattern: (arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})

Required: Yes

versionToken

The version token of the table.

Length Constraints: Minimum length of 1. Maximum length of 2048.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteTableBucket

Service: Amazon S3 Tables

Deletes a table bucket. For more information, see [Deleting a table bucket](#) in the *Amazon Simple Storage Service User Guide*.

Permissions

You must have the `s3tables:DeleteTableBucket` permission to use this operation.

Request Syntax

```
DELETE /buckets/tableBucketARN HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket.

Pattern: (`arn:aws``[-a-z0-9]*``:[a-z0-9]+``:[-a-z0-9]*``:[0-9]{12}``:bucket/``[a-z0-9_-]{3,63}`)

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)

- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteTableBucketEncryption

Service: Amazon S3 Tables

Deletes the encryption configuration for a table bucket.

Permissions

You must have the `s3tables:DeleteTableBucketEncryption` permission to use this operation.

Request Syntax

```
DELETE /buckets/tableBucketARN/encryption HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket.

Pattern: (`arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63}`)

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)

- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteTableBucketMetricsConfiguration

Service: Amazon S3 Tables

Deletes the metrics configuration for a table bucket.

Permissions

You must have the `s3tables:DeleteTableBucketMetricsConfiguration` permission to use this operation.

Request Syntax

```
DELETE /buckets/tableBucketARN/metrics HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket.

Pattern: (`arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63}`)

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)

- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteTableBucketPolicy

Service: Amazon S3 Tables

Deletes a table bucket policy. For more information, see [Deleting a table bucket policy](#) in the *Amazon Simple Storage Service User Guide*.

Permissions

You must have the `s3tables:DeleteTableBucketPolicy` permission to use this operation.

Request Syntax

```
DELETE /buckets/tableBucketARN/policy HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket.

Pattern: (`arn:aws``[-a-z0-9]*``:[a-z0-9]+``:[-a-z0-9]*``:[0-9]{12}``:bucket/``[a-z0-9_-]{3,63}`)

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)

- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteTableBucketReplication

Service: Amazon S3 Tables

Deletes the replication configuration for a table bucket. After deletion, new table updates will no longer be replicated to destination buckets, though existing replicated tables will remain in destination buckets.

Permissions

You must have the `s3tables:DeleteTableBucketReplication` permission to use this operation.

Request Syntax

```
DELETE /table-bucket-replication?  
tableBucketARN=tableBucketARN&versionToken=versionToken HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket.

Pattern: (`arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63}`)

Required: Yes

versionToken

A version token from a previous `GetTableBucketReplication` call. Use this token to ensure you're deleting the expected version of the configuration.

Length Constraints: Minimum length of 1. Maximum length of 2048.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

AccessDeniedException

The action cannot be performed because you do not have the required permission.

HTTP Status Code: 403

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteTablePolicy

Service: Amazon S3 Tables

Deletes a table policy. For more information, see [Deleting a table policy](#) in the *Amazon Simple Storage Service User Guide*.

Permissions

You must have the `s3tables:DeleteTablePolicy` permission to use this operation.

Request Syntax

```
DELETE /tables/tableBucketARN/namespace/name/policy HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

name

The table name.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

namespace

The namespace associated with the table.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket that contains the table.

Pattern: (arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteTableReplication

Service: Amazon S3 Tables

Deletes the replication configuration for a specific table. After deletion, new updates to this table will no longer be replicated to destination tables, though existing replicated copies will remain in destination buckets.

Permissions

You must have the `s3tables:DeleteTableReplication` permission to use this operation.

Request Syntax

```
DELETE /table-replication?tableArn=tableArn&versionToken=versionToken HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

tableArn

The Amazon Resource Name (ARN) of the table.

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: (`arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63}/table/[a-zA-Z0-9-_{1,255})`

Required: Yes

versionToken

A version token from a previous `GetTableReplication` call. Use this token to ensure you're deleting the expected version of the configuration.

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

AccessDeniedException

The action cannot be performed because you do not have the required permission.

HTTP Status Code: 403

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetNamespace

Service: Amazon S3 Tables

Gets details about a namespace. For more information, see [Table namespaces](#) in the *Amazon Simple Storage Service User Guide*.

Permissions

You must have the `s3tables:GetNamespace` permission to use this operation.

Request Syntax

```
GET /namespaces/tableBucketARN/namespace HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

[namespace](#)

The name of the namespace.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

[tableBucketARN](#)

The Amazon Resource Name (ARN) of the table bucket.

Pattern: `(arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "createdAt": "string",
  "createdBy": "string",
  "namespace": [ "string" ],
  "namespaceId": "string",
  "ownerAccountId": "string",
  "tableBucketId": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[createdAt](#)

The date and time the namespace was created at.

Type: Timestamp

[createdBy](#)

The ID of the account that created the namespace.

Type: String

Length Constraints: Fixed length of 12.

Pattern: `[0-9].*`

[namespace](#)

The name of the namespace.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

namespaceId

The unique identifier of the namespace.

Type: String

ownerAccountId

The ID of the account that owns the namespace.

Type: String

Length Constraints: Fixed length of 12.

Pattern: [0-9].*

tableBucketId

The unique identifier of the table bucket containing this namespace.

Type: String

Errors

AccessDeniedException

The action cannot be performed because you do not have the required permission.

HTTP Status Code: 403

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetTable

Service: Amazon S3 Tables

Gets details about a table. For more information, see [S3 Tables](#) in the *Amazon Simple Storage Service User Guide*.

Permissions

You must have the `s3tables:GetTable` permission to use this operation.

Request Syntax

```
GET /get-table?  
name=name&namespace=namespace&tableArn=tableArn&tableBucketARN=tableBucketARN HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the table.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

namespace

The name of the namespace the table is associated with.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

tableArn

The Amazon Resource Name (ARN) of the table.

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `(arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63}/table/[a-zA-Z0-9_-]{1,255})`

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket associated with the table.

Pattern: (arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "createdAt": "string",
  "createdBy": "string",
  "format": "string",
  "managedByService": "string",
  "managedTableInformation": {
    "replicationInformation": {
      "sourceTableARN": "string"
    }
  },
  "metadataLocation": "string",
  "modifiedAt": "string",
  "modifiedBy": "string",
  "name": "string",
  "namespace": [ "string" ],
  "namespaceId": "string",
  "ownerAccountId": "string",
  "tableARN": "string",
  "tableBucketId": "string",
  "type": "string",
  "versionToken": "string",
  "warehouseLocation": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

createdAt

The date and time the table bucket was created at.

Type: Timestamp

createdBy

The ID of the account that created the table.

Type: String

Length Constraints: Fixed length of 12.

Pattern: `[0-9].*`

format

The format of the table.

Type: String

Valid Values: ICEBERG

managedByService

The service that manages the table.

Type: String

managedTableInformation

If this table is managed by S3 Tables, contains additional information such as replication details.

Type: [ManagedTableInformation](#) object

metadataLocation

The metadata location of the table.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

modifiedAt

The date and time the table was last modified on.

Type: Timestamp

modifiedBy

The ID of the account that last modified the table.

Type: String

Length Constraints: Fixed length of 12.

Pattern: [0-9].*

name

The name of the table.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: [0-9a-z_]*

namespace

The namespace associated with the table.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: [0-9a-z_]*

namespaceId

The unique identifier of the namespace containing this table.

Type: String

ownerAccountId

The ID of the account that owns the table.

Type: String

Length Constraints: Fixed length of 12.

Pattern: `[0-9].*`

tableARN

The Amazon Resource Name (ARN) of the table.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `(arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63}/table/[a-zA-Z0-9_-]{1,255})`

tableBucketId

The unique identifier of the table bucket containing this table.

Type: String

type

The type of the table.

Type: String

Valid Values: `customer` | `aws`

versionToken

The version token of the table.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

warehouseLocation

The warehouse location of the table.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Errors

AccessDeniedException

The action cannot be performed because you do not have the required permission.

HTTP Status Code: 403

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetTableBucket

Service: Amazon S3 Tables

Gets details on a table bucket. For more information, see [Viewing details about an Amazon S3 table bucket](#) in the *Amazon Simple Storage Service User Guide*.

Permissions

You must have the `s3tables:GetTableBucket` permission to use this operation.

Request Syntax

```
GET /buckets/tableBucketARN HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket.

Pattern: `(arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "arn": "string",
  "createdAt": "string",
  "name": "string",
```

```
"ownerAccountId": "string",  
"tableBucketId": "string",  
"type": "string"  
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

arn

The Amazon Resource Name (ARN) of the table bucket.

Type: String

Pattern: (arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})

createdAt

The date and time the table bucket was created.

Type: Timestamp

name

The name of the table bucket

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Pattern: [0-9a-z-]*

ownerAccountId

The ID of the account that owns the table bucket.

Type: String

Length Constraints: Fixed length of 12.

Pattern: [0-9].*

tableBucketId

The unique identifier of the table bucket.

Type: String

type

The type of the table bucket.

Type: String

Valid Values: customer | aws

Errors

AccessDeniedException

The action cannot be performed because you do not have the required permission.

HTTP Status Code: 403

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetTableBucketEncryption

Service: Amazon S3 Tables

Gets the encryption configuration for a table bucket.

Permissions

You must have the `s3tables:GetTableBucketEncryption` permission to use this operation.

Request Syntax

```
GET /buckets/tableBucketARN/encryption HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

[tableBucketARN](#)

The Amazon Resource Name (ARN) of the table bucket.

Pattern: (`arn:aws``[-a-z0-9]*``:[a-z0-9]+``:[-a-z0-9]*``:[0-9]{12}``:bucket/``[a-z0-9_-]{3,63}`)

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "encryptionConfiguration": {
    "kmsKeyArn": "string",
    "sseAlgorithm": "string"
```

```
}  
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[encryptionConfiguration](#)

The encryption configuration for the table bucket.

Type: [EncryptionConfiguration](#) object

Errors

AccessDeniedException

The action cannot be performed because you do not have the required permission.

HTTP Status Code: 403

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetTableBucketMaintenanceConfiguration

Service: Amazon S3 Tables

Gets details about a maintenance configuration for a given table bucket. For more information, see [Amazon S3 table bucket maintenance](#) in the *Amazon Simple Storage Service User Guide*.

Permissions

You must have the `s3tables:GetTableBucketMaintenanceConfiguration` permission to use this operation.

Request Syntax

```
GET /buckets/tableBucketARN/maintenance HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket associated with the maintenance configuration.

Pattern: `(arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
```

```
"configuration": {
  "string" : {
    "settings": { ... },
    "status": "string"
  }
},
"tableBucketARN": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

configuration

Details about the maintenance configuration for the table bucket.

Type: String to [TableBucketMaintenanceConfigurationValue](#) object map

Valid Keys: icebergUnreferencedFileRemoval

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket associated with the maintenance configuration.

Type: String

Pattern: (arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetTableBucketMetricsConfiguration

Service: Amazon S3 Tables

Gets the metrics configuration for a table bucket.

Permissions

You must have the `s3tables:GetTableBucketMetricsConfiguration` permission to use this operation.

Request Syntax

```
GET /buckets/tableBucketARN/metrics HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket.

Pattern: (`arn:aws``[-a-z0-9]*``:[a-z0-9]+``:[-a-z0-9]*``:[0-9]{12}``:bucket/``[a-z0-9_-]{3,63}`)

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "id": "string",
  "tableBucketARN": "string"
}
```


Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

id

The unique identifier of the metrics configuration.

Type: String

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket.

Type: String

Pattern: (arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetTableBucketPolicy

Service: Amazon S3 Tables

Gets details about a table bucket policy. For more information, see [Viewing a table bucket policy](#) in the *Amazon Simple Storage Service User Guide*.

Permissions

You must have the `s3tables:GetTableBucketPolicy` permission to use this operation.

Request Syntax

```
GET /buckets/tableBucketARN/policy HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket.

Pattern: (`arn:aws``[-a-z0-9]*``:[a-z0-9]+``:[-a-z0-9]*``:[0-9]{12}``:bucket/``[a-z0-9_-]{3,63}`)

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "resourcePolicy": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

resourcePolicy

The JSON that defines the policy.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 20480.

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetTableBucketReplication

Service: Amazon S3 Tables

Retrieves the replication configuration for a table bucket. This operation returns the IAM role, versionToken, and replication rules that define how tables in this bucket are replicated to other buckets.

Permissions

You must have the `s3tables:GetTableBucketReplication` permission to use this operation.

Request Syntax

```
GET /table-bucket-replication?tableBucketARN=tableBucketARN HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket.

Pattern: `(arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
```

```
"configuration": {
  "role": "string",
  "rules": [
    {
      "destinations": [
        {
          "destinationTableBucketARN": "string"
        }
      ]
    }
  ],
  "versionToken": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

configuration

The replication configuration for the table bucket, including the IAM role and replication rules.

Type: [TableBucketReplicationConfiguration](#) object

versionToken

A version token that represents the current state of the replication configuration. Use this token when updating the configuration to ensure consistency.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Errors

AccessDeniedException

The action cannot be performed because you do not have the required permission.

HTTP Status Code: 403

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetTableBucketStorageClass

Service: Amazon S3 Tables

Retrieves the storage class configuration for a specific table. This allows you to view the storage class settings that apply to an individual table, which may differ from the table bucket's default configuration.

Permissions

You must have the `s3tables:GetTableBucketStorageClass` permission to use this operation.

Request Syntax

```
GET /buckets/tableBucketARN/storage-class HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket.

Pattern: (`arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63}`)

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "storageClassConfiguration": {
```

```
    "storageClass": "string"  
  }  
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

storageClassConfiguration

The storage class configuration for the table bucket.

Type: [StorageClassConfiguration](#) object

Errors

AccessDeniedException

The action cannot be performed because you do not have the required permission.

HTTP Status Code: 403

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetTableEncryption

Service: Amazon S3 Tables

Gets the encryption configuration for a table.

Permissions

You must have the `s3tables:GetTableEncryption` permission to use this operation.

Request Syntax

```
GET /tables/tableBucketARN/namespace/name/encryption HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the table.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

namespace

The namespace associated with the table.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket containing the table.

Pattern: `(arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "encryptionConfiguration": {
    "kmsKeyArn": "string",
    "sseAlgorithm": "string"
  }
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

encryptionConfiguration

The encryption configuration for the table.

Type: [EncryptionConfiguration](#) object

Errors

AccessDeniedException

The action cannot be performed because you do not have the required permission.

HTTP Status Code: 403

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetTableMaintenanceConfiguration

Service: Amazon S3 Tables

Gets details about the maintenance configuration of a table. For more information, see [S3 Tables maintenance](#) in the *Amazon Simple Storage Service User Guide*.

Permissions

- You must have the `s3tables:GetTableMaintenanceConfiguration` permission to use this operation.
- You must have the `s3tables:GetTableData` permission to use set the compaction strategy to sort or zorder.

Request Syntax

```
GET /tables/tableBucketARN/namespace/name/maintenance HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

[name](#)

The name of the table.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

[namespace](#)

The namespace associated with the table.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket.

Pattern: (arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "configuration": {
    "string" : {
      "settings": { ... },
      "status": "string"
    }
  },
  "tableARN": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

configuration

Details about the maintenance configuration for the table bucket.

Type: String to [TableMaintenanceConfigurationValue](#) object map

Valid Keys: icebergCompaction | icebergSnapshotManagement

tableARN

The Amazon Resource Name (ARN) of the table.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: (`arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63}/table/[a-zA-Z0-9-]{1,255}`)

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetTableMaintenanceJobStatus

Service: Amazon S3 Tables

Gets the status of a maintenance job for a table. For more information, see [S3 Tables maintenance](#) in the *Amazon Simple Storage Service User Guide*.

Permissions

You must have the `s3tables:GetTableMaintenanceJobStatus` permission to use this operation.

Request Syntax

```
GET /tables/tableBucketARN/namespace/name/maintenance-job-status HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the table containing the maintenance job status you want to check.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

namespace

The name of the namespace the table is associated with.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket.

Pattern: (arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "status": {
    "string" : {
      "failureMessage": "string",
      "lastRunTimestamp": "string",
      "status": "string"
    }
  },
  "tableARN": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

status

The status of the maintenance job.

Type: String to [TableMaintenanceJobStatusValue](#) object map

Valid Keys: icebergCompaction | icebergSnapshotManagement |
icebergUnreferencedFileRemoval

tableARN

The Amazon Resource Name (ARN) of the table.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: (arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63}/table/[a-zA-Z0-9_-]{1,255})

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetTableMetadataLocation

Service: Amazon S3 Tables

Gets the location of the table metadata.

Permissions

You must have the `s3tables:GetTableMetadataLocation` permission to use this operation.

Request Syntax

```
GET /tables/tableBucketARN/namespace/name/metadata-location HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the table.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

namespace

The namespace of the table.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket.

Pattern: (arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "metadataLocation": "string",
  "versionToken": "string",
  "warehouseLocation": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

metadataLocation

The metadata location.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

versionToken

The version token.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

warehouseLocation

The warehouse location.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetTablePolicy

Service: Amazon S3 Tables

Gets details about a table policy. For more information, see [Viewing a table policy](#) in the *Amazon Simple Storage Service User Guide*.

Permissions

You must have the `s3tables:GetTablePolicy` permission to use this operation.

Request Syntax

```
GET /tables/tableBucketARN/namespace/name/policy HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the table.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

namespace

The namespace associated with the table.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket that contains the table.

Pattern: `(arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "resourcePolicy": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[resourcePolicy](#)

The JSON that defines the policy.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 20480.

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetTableRecordExpirationConfiguration

Service: Amazon S3 Tables

Retrieves the expiration configuration settings for records in a table, and the status of the configuration. If the status of the configuration is enabled, records expire and are automatically removed from the table after the specified number of days.

Permissions

You must have the `s3tables:GetTableRecordExpirationConfiguration` permission to use this operation.

Request Syntax

```
GET /table-record-expiration?tableArn=tableArn HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

tableArn

The Amazon Resource Name (ARN) of the table.

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `(arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63}/table/[a-zA-Z0-9-]{1,255})`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json
```



```
{
  "configuration": {
    "settings": {
      "days": number
    },
    "status": "string"
  }
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

configuration

The record expiration configuration for the table, including the status and retention settings.

Type: [TableRecordExpirationConfigurationValue](#) object

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

MethodNotAllowedException

The requested operation is not allowed on this resource. This may occur when attempting to modify a resource that is managed by a service or has restrictions that prevent the operation.

HTTP Status Code: 405

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetTableRecordExpirationJobStatus

Service: Amazon S3 Tables

Retrieves the status, metrics, and details of the latest record expiration job for a table. This includes when the job ran, and whether it succeeded or failed. If the job ran successfully, this also includes statistics about the records that were removed.

Permissions

You must have the `s3tables:GetTableRecordExpirationJobStatus` permission to use this operation.

Request Syntax

```
GET /table-record-expiration-job-status?tableArn=tableArn HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

tableArn

The Amazon Resource Name (ARN) of the table.

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `(arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63}/table/[a-zA-Z0-9_-]{1,255})`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json
```

```
{
  "failureMessage": "string",
  "lastRunTimestamp": "string",
  "metrics": {
    "deletedDataFiles": number,
    "deletedRecords": number,
    "removedFilesSize": number
  },
  "status": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

failureMessage

If the job failed, this field contains an error message describing the failure reason.

Type: String

lastRunTimestamp

The timestamp when the expiration job was last executed.

Type: Timestamp

metrics

Metrics about the most recent expiration job execution, including the number of records and files deleted.

Type: [TableRecordExpirationJobMetrics](#) object

status

The current status of the most recent expiration job.

Type: String

Valid Values: NotYetRun | Successful | Failed | Disabled

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

MethodNotAllowedException

The requested operation is not allowed on this resource. This may occur when attempting to modify a resource that is managed by a service or has restrictions that prevent the operation.

HTTP Status Code: 405

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)

- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetTableReplication

Service: Amazon S3 Tables

Retrieves the replication configuration for a specific table.

Permissions

You must have the `s3tables:GetTableReplication` permission to use this operation.

Request Syntax

```
GET /table-replication?tableArn=tableArn HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

tableArn

The Amazon Resource Name (ARN) of the table.

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: (arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63}/table/[a-zA-Z0-9-_{1,255})

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "configuration": {
    "role": "string",
```

```
  "rules": [
    {
      "destinations": [
        {
          "destinationTableBucketARN": "string"
        }
      ]
    }
  ],
  "versionToken": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

configuration

The replication configuration for the table, including the IAM role and replication rules.

Type: [TableReplicationConfiguration](#) object

versionToken

A version token that represents the current state of the table's replication configuration. Use this token when updating the configuration to ensure consistency.

Type: String

Errors

AccessDeniedException

The action cannot be performed because you do not have the required permission.

HTTP Status Code: 403

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)

- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetTableReplicationStatus

Service: Amazon S3 Tables

Retrieves the replication status for a table, including the status of replication to each destination. This operation provides visibility into replication health and progress.

Permissions

You must have the `s3tables:GetTableReplicationStatus` permission to use this operation.

Request Syntax

```
GET /replication-status?tableArn=tableArn HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

tableArn

The Amazon Resource Name (ARN) of the table.

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `(arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63}/table/[a-zA-Z0-9-]{1,255})`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json
```

```
{
  "destinations": [
    {
      "destinationTableArn": "string",
      "destinationTableBucketArn": "string",
      "failureMessage": "string",
      "lastSuccessfulReplicatedUpdate": {
        "metadataLocation": "string",
        "timestamp": "string"
      },
      "replicationStatus": "string"
    }
  ],
  "sourceTableArn": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[destinations](#)

An array of status information for each replication destination, including the current state, last successful update, and any error messages.

Type: Array of [ReplicationDestinationStatusModel](#) objects

Array Members: Minimum number of 1 item. Maximum number of 5 items.

[sourceTableArn](#)

The Amazon Resource Name (ARN) of the source table being replicated.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: (arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63}/table/[a-zA-Z0-9_-]{1,255})

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)

- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetTableStorageClass

Service: Amazon S3 Tables

Retrieves the storage class configuration for a specific table. This allows you to view the storage class settings that apply to an individual table, which may differ from the table bucket's default configuration.

Permissions

You must have the `s3tables:GetTableStorageClass` permission to use this operation.

Request Syntax

```
GET /tables/tableBucketARN/namespace/name/storage-class HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the table.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

namespace

The namespace associated with the table.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket that contains the table.

Pattern: (arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "storageClassConfiguration": {
    "storageClass": "string"
  }
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

storageClassConfiguration

The storage class configuration for the table.

Type: [StorageClassConfiguration](#) object

Errors

AccessDeniedException

The action cannot be performed because you do not have the required permission.

HTTP Status Code: 403

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListNamespaces

Service: Amazon S3 Tables

Lists the namespaces within a table bucket. For more information, see [Table namespaces](#) in the *Amazon Simple Storage Service User Guide*.

Permissions

You must have the `s3tables:ListNamespaces` permission to use this operation.

Request Syntax

```
GET /namespaces/tableBucketARN?  
continuationToken=continuationToken&maxNamespaces=maxNamespaces&prefix=prefix HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

[continuationToken](#)

ContinuationToken indicates to Amazon S3 that the list is being continued on this bucket with a token. ContinuationToken is obfuscated and is not a real key. You can use this ContinuationToken for pagination of the list results.

Length Constraints: Minimum length of 1. Maximum length of 2048.

[maxNamespaces](#)

The maximum number of namespaces to return in the list.

Valid Range: Minimum value of 1. Maximum value of 1000.

[prefix](#)

The prefix of the namespaces.

Length Constraints: Minimum length of 1. Maximum length of 255.

[tableBucketARN](#)

The Amazon Resource Name (ARN) of the table bucket.

Pattern: (arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "continuationToken": "string",
  "namespaces": [
    {
      "createdAt": "string",
      "createdBy": "string",
      "namespace": [ "string" ],
      "namespaceId": "string",
      "ownerAccountId": "string",
      "tableBucketId": "string"
    }
  ]
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

continuationToken

The ContinuationToken for pagination of the list results.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

[namespaces](#)

A list of namespaces.

Type: Array of [NamespaceSummary](#) objects

Errors

AccessDeniedException

The action cannot be performed because you do not have the required permission.

HTTP Status Code: 403

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListTableBuckets

Service: Amazon S3 Tables

Lists table buckets for your account. For more information, see [S3 Table buckets](#) in the *Amazon Simple Storage Service User Guide*.

Permissions

You must have the `s3tables:ListTableBuckets` permission to use this operation.

Request Syntax

```
GET /buckets?
continuationToken=continuationToken&maxBuckets=maxBuckets&prefix=prefix&type=type
HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

[continuationToken](#)

ContinuationToken indicates to Amazon S3 that the list is being continued on this bucket with a token. ContinuationToken is obfuscated and is not a real key. You can use this ContinuationToken for pagination of the list results.

Length Constraints: Minimum length of 1. Maximum length of 2048.

[maxBuckets](#)

The maximum number of table buckets to return in the list.

Valid Range: Minimum value of 1. Maximum value of 1000.

[prefix](#)

The prefix of the table buckets.

Length Constraints: Minimum length of 1. Maximum length of 63.

[type](#)

The type of table buckets to filter by in the list.

Valid Values: customer | aws

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "continuationToken": "string",
  "tableBuckets": [
    {
      "arn": "string",
      "createdAt": "string",
      "name": "string",
      "ownerAccountId": "string",
      "tableBucketId": "string",
      "type": "string"
    }
  ]
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

continuationToken

You can use this ContinuationToken for pagination of the list results.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

tableBuckets

A list of table buckets.

Type: Array of [TableBucketSummary](#) objects

Errors

AccessDeniedException

The action cannot be performed because you do not have the required permission.

HTTP Status Code: 403

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListTables

Service: Amazon S3 Tables

List tables in the given table bucket. For more information, see [S3 Tables](#) in the *Amazon Simple Storage Service User Guide*.

Permissions

You must have the `s3tables:ListTables` permission to use this operation.

Request Syntax

```
GET /tables/tableBucketARN?  
continuationToken=continuationToken&maxTables=maxTables&namespace=namespace&prefix=prefix  
HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

[continuationToken](#)

ContinuationToken indicates to Amazon S3 that the list is being continued on this bucket with a token. ContinuationToken is obfuscated and is not a real key. You can use this ContinuationToken for pagination of the list results.

Length Constraints: Minimum length of 1. Maximum length of 2048.

[maxTables](#)

The maximum number of tables to return.

Valid Range: Minimum value of 1. Maximum value of 1000.

[namespace](#)

The namespace of the tables.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

prefix

The prefix of the tables.

Length Constraints: Minimum length of 1. Maximum length of 255.

tableBucketARN

The Amazon resource Name (ARN) of the table bucket.

Pattern: (arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "continuationToken": "string",
  "tables": [
    {
      "createdAt": "string",
      "managedByService": "string",
      "modifiedAt": "string",
      "name": "string",
      "namespace": [ "string" ],
      "namespaceId": "string",
      "tableARN": "string",
      "tableBucketId": "string",
      "type": "string"
    }
  ]
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

continuationToken

You can use this ContinuationToken for pagination of the list results.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

tables

A list of tables.

Type: Array of [TableSummary](#) objects

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListTagsForResource

Service: Amazon S3 Tables

Lists all of the tags applied to a specified Amazon S3 Tables resource. Each tag is a label consisting of a key and value pair. Tags can help you organize, track costs for, and control access to resources.

Note

For a list of S3 resources that support tagging, see [Managing tags for Amazon S3 resources](#).

Permissions

For tables and table buckets, you must have the `s3tables:ListTagsForResource` permission to use this operation.

Request Syntax

```
GET /tag/resourceArn HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

resourceArn

The Amazon Resource Name (ARN) of the Amazon S3 Tables resource that you want to list tags for. The tagged resource can be a table bucket or a table. For a list of all S3 resources that support tagging, see [Managing tags for Amazon S3 resources](#).

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/.+`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "tags": {
    "string" : "string"
  }
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

tags

The user-defined tags that are applied to the resource. For more information, see [Tagging for cost allocation or attribute-based access control \(ABAC\)](#).

Type: String to string map

Key Length Constraints: Minimum length of 1. Maximum length of 128.

Value Length Constraints: Minimum length of 0. Maximum length of 256.

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutTableBucketEncryption

Service: Amazon S3 Tables

Sets the encryption configuration for a table bucket.

Permissions

You must have the `s3tables:PutTableBucketEncryption` permission to use this operation.

Note

If you choose SSE-KMS encryption you must grant the S3 Tables maintenance principal access to your KMS key. For more information, see [Permissions requirements for S3 Tables SSE-KMS encryption](#) in the *Amazon Simple Storage Service User Guide*.

Request Syntax

```
PUT /buckets/tableBucketARN/encryption HTTP/1.1
Content-type: application/json
```

```
{
  "encryptionConfiguration": {
    "kmsKeyArn": "string",
    "sseAlgorithm": "string"
  }
}
```

URI Request Parameters

The request uses the following URI parameters.

[tableBucketARN](#)

The Amazon Resource Name (ARN) of the table bucket.

Pattern: `(arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})`

Required: Yes

Request Body

The request accepts the following data in JSON format.

[encryptionConfiguration](#)

The encryption configuration to apply to the table bucket.

Type: [EncryptionConfiguration](#) object

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutTableBucketMaintenanceConfiguration

Service: Amazon S3 Tables

Creates a new maintenance configuration or replaces an existing maintenance configuration for a table bucket. For more information, see [Amazon S3 table bucket maintenance](#) in the *Amazon Simple Storage Service User Guide*.

Permissions

You must have the `s3tables:PutTableBucketMaintenanceConfiguration` permission to use this operation.

Request Syntax

```
PUT /buckets/tableBucketARN/maintenance/type HTTP/1.1
Content-type: application/json

{
  "value": {
    "settings": { ... },
    "status": "string"
  }
}
```

URI Request Parameters

The request uses the following URI parameters.

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket associated with the maintenance configuration.

Pattern: (`arn:aws``[-a-z0-9]*``:[a-z0-9]+``:[-a-z0-9]*``:[0-9]{12}``:bucket/``[a-z0-9_-]{3,63}`)

Required: Yes

type

The type of the maintenance configuration.

Valid Values: `icebergUnreferencedFileRemoval`

Required: Yes

Request Body

The request accepts the following data in JSON format.

value

Defines the values of the maintenance configuration for the table bucket.

Type: [TableBucketMaintenanceConfigurationValue](#) object

Required: Yes

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutTableBucketMetricsConfiguration

Service: Amazon S3 Tables

Sets the metrics configuration for a table bucket.

Permissions

You must have the `s3tables:PutTableBucketMetricsConfiguration` permission to use this operation.

Request Syntax

```
PUT /buckets/tableBucketARN/metrics HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket.

Pattern: (`arn:aws``[-a-z0-9]*``:[a-z0-9]+``:[-a-z0-9]*``:[0-9]{12}``:bucket/``[a-z0-9_-]{3,63}`)

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)

- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutTableBucketPolicy

Service: Amazon S3 Tables

Creates a new table bucket policy or replaces an existing table bucket policy for a table bucket. For more information, see [Adding a table bucket policy](#) in the *Amazon Simple Storage Service User Guide*.

Permissions

You must have the `s3tables:PutTableBucketPolicy` permission to use this operation.

Request Syntax

```
PUT /buckets/tableBucketARN/policy HTTP/1.1
Content-type: application/json

{
  "resourcePolicy": "string"
}
```

URI Request Parameters

The request uses the following URI parameters.

[tableBucketARN](#)

The Amazon Resource Name (ARN) of the table bucket.

Pattern: (`arn:aws``[-a-z0-9]*``:[a-z0-9]+``:[-a-z0-9]*``:[0-9]{12}``:bucket/``[a-z0-9_-]{3,63}`)

Required: Yes

Request Body

The request accepts the following data in JSON format.

[resourcePolicy](#)

The JSON that defines the policy.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 20480.

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutTableBucketReplication

Service: Amazon S3 Tables

Creates or updates the replication configuration for a table bucket. This operation defines how tables in the source bucket are replicated to destination buckets. Replication helps ensure data availability and disaster recovery across regions or accounts.

Permissions

- You must have the `s3tables:PutTableBucketReplication` permission to use this operation. The IAM role specified in the configuration must have permissions to read from the source bucket and write permissions to all destination buckets.
- You must also have the following permissions:
 - `s3tables:GetTable` permission on the source table.
 - `s3tables:ListTables` permission on the bucket containing the table.
 - `s3tables:CreateTable` permission for the destination.
 - `s3tables:CreateNamespace` permission for the destination.
 - `s3tables:GetTableMaintenanceConfig` permission for the source bucket.
 - `s3tables:PutTableMaintenanceConfig` permission for the destination bucket.
- You must have `iam:PassRole` permission with condition allowing roles to be passed to `replication.s3tables.amazonaws.com`.

Request Syntax

```
PUT /table-bucket-replication?tableBucketARN=tableBucketARN&versionToken=versionToken
HTTP/1.1
Content-type: application/json

{
  "configuration": {
    "role": "string",
    "rules": [
      {
        "destinations": [
          {
            "destinationTableBucketARN": "string"
          }
        ]
      }
    ]
  }
}
```

```
}  
  ]  
}  
}
```

URI Request Parameters

The request uses the following URI parameters.

tableBucketARN

The Amazon Resource Name (ARN) of the source table bucket.

Pattern: (arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})

Required: Yes

versionToken

A version token from a previous GetTableBucketReplication call. Use this token to ensure you're updating the expected version of the configuration.

Length Constraints: Minimum length of 1. Maximum length of 2048.

Request Body

The request accepts the following data in JSON format.

configuration

The replication configuration to apply, including the IAM role and replication rules.

Type: [TableBucketReplicationConfiguration](#) object

Required: Yes

Response Syntax

```
HTTP/1.1 200  
Content-type: application/json
```

```
{
  "status": "string",
  "versionToken": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

status

The status of the replication configuration operation.

Type: String

versionToken

A new version token representing the updated replication configuration.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Errors

AccessDeniedException

The action cannot be performed because you do not have the required permission.

HTTP Status Code: 403

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutTableBucketStorageClass

Service: Amazon S3 Tables

Sets or updates the storage class configuration for a table bucket. This configuration serves as the default storage class for all new tables created in the bucket, allowing you to optimize storage costs at the bucket level.

Permissions

You must have the `s3tables:PutTableBucketStorageClass` permission to use this operation.

Request Syntax

```
PUT /buckets/tableBucketARN/storage-class HTTP/1.1
Content-type: application/json

{
  "storageClassConfiguration": {
    "storageClass": "string"
  }
}
```

URI Request Parameters

The request uses the following URI parameters.

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket.

Pattern: (`arn:aws``[-a-z0-9]*``:[a-z0-9]+``:[-a-z0-9]*``:[0-9]{12}``:bucket/``[a-z0-9_-]{3,63}`)

Required: Yes

Request Body

The request accepts the following data in JSON format.

storageClassConfiguration

The storage class configuration to apply to the table bucket. This configuration will serve as the default for new tables created in this bucket.

Type: [StorageClassConfiguration](#) object

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutTableMaintenanceConfiguration

Service: Amazon S3 Tables

Creates a new maintenance configuration or replaces an existing maintenance configuration for a table. For more information, see [S3 Tables maintenance](#) in the *Amazon Simple Storage Service User Guide*.

Permissions

You must have the `s3tables:PutTableMaintenanceConfiguration` permission to use this operation.

Request Syntax

```
PUT /tables/tableBucketARN/namespace/name/maintenance/type HTTP/1.1
Content-type: application/json

{
  "value": {
    "settings": { ... },
    "status": "string"
  }
}
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the table.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

namespace

The namespace of the table.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

[tableBucketARN](#)

The Amazon Resource Name (ARN) of the table associated with the maintenance configuration.

Pattern: `(arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})`

Required: Yes

[type](#)

The type of the maintenance configuration.

Valid Values: `icebergCompaction` | `icebergSnapshotManagement`

Required: Yes

Request Body

The request accepts the following data in JSON format.

[value](#)

Defines the values of the maintenance configuration for the table.

Type: [TableMaintenanceConfigurationValue](#) object

Required: Yes

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)

- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutTablePolicy

Service: Amazon S3 Tables

Creates a new table policy or replaces an existing table policy for a table. For more information, see [Adding a table policy](#) in the *Amazon Simple Storage Service User Guide*.

Permissions

You must have the `s3tables:PutTablePolicy` permission to use this operation.

Request Syntax

```
PUT /tables/tableBucketARN/namespace/name/policy HTTP/1.1
Content-type: application/json

{
  "resourcePolicy": "string"
}
```

URI Request Parameters

The request uses the following URI parameters.

[name](#)

The name of the table.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

[namespace](#)

The namespace associated with the table.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket that contains the table.

Pattern: (`arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63}`)

Required: Yes

Request Body

The request accepts the following data in JSON format.

resourcePolicy

The JSON that defines the policy.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 20480.

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutTableRecordExpirationConfiguration

Service: Amazon S3 Tables

Creates or updates the expiration configuration settings for records in a table, including the status of the configuration. If you enable record expiration for a table, records expire and are automatically removed from the table after the number of days that you specify.

Permissions

You must have the `s3tables:PutTableRecordExpirationConfiguration` permission to use this operation.

Request Syntax

```
PUT /table-record-expiration?tableArn=tableArn HTTP/1.1
Content-type: application/json
```

```
{
  "value": {
    "settings": {
      "days": number
    },
    "status": "string"
  }
}
```

URI Request Parameters

The request uses the following URI parameters.

tableArn

The Amazon Resource Name (ARN) of the table.

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: (arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63}/table/[a-zA-Z0-9-]{1,255})

Required: Yes

Request Body

The request accepts the following data in JSON format.

value

The record expiration configuration to apply to the table, including the status (enabled or disabled) and retention period in days.

Type: [TableRecordExpirationConfigurationValue](#) object

Required: Yes

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

MethodNotAllowedException

The requested operation is not allowed on this resource. This may occur when attempting to modify a resource that is managed by a service or has restrictions that prevent the operation.

HTTP Status Code: 405

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutTableReplication

Service: Amazon S3 Tables

Creates or updates the replication configuration for a specific table. This operation allows you to define table-level replication independently of bucket-level replication, providing granular control over which tables are replicated and where.

Permissions

- You must have the `s3tables:PutTableReplication` permission to use this operation. The IAM role specified in the configuration must have permissions to read from the source table and write to all destination tables.
- You must also have the following permissions:
 - `s3tables:GetTable` permission on the source table being replicated.
 - `s3tables:CreateTable` permission for the destination.
 - `s3tables:CreateNamespace` permission for the destination.
 - `s3tables:GetTableMaintenanceConfig` permission for the source table.
 - `s3tables:PutTableMaintenanceConfig` permission for the destination table.
- You must have `iam:PassRole` permission with condition allowing roles to be passed to `replication.s3tables.amazonaws.com`.

Request Syntax

```
PUT /table-replication?tableArn=tableArn&versionToken=versionToken HTTP/1.1
Content-type: application/json
```

```
{
  "configuration": {
    "role": "string",
    "rules": [
      {
        "destinations": [
          {
            "destinationTableBucketARN": "string"
          }
        ]
      }
    ]
  }
}
```



```
}  
}
```

URI Request Parameters

The request uses the following URI parameters.

[tableArn](#)

The Amazon Resource Name (ARN) of the source table.

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: (arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63}/table/[a-zA-Z0-9-_{1,255})

Required: Yes

[versionToken](#)

A version token from a previous GetTableReplication call. Use this token to ensure you're updating the expected version of the configuration.

Request Body

The request accepts the following data in JSON format.

[configuration](#)

The replication configuration to apply to the table, including the IAM role and replication rules.

Type: [TableReplicationConfiguration](#) object

Required: Yes

Response Syntax

```
HTTP/1.1 200  
Content-type: application/json  
  
{
```

```
"status": "string",  
"versionToken": "string"  
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[status](#)

The status of the replication configuration operation.

Type: String

[versionToken](#)

A new version token representing the updated replication configuration.

Type: String

Errors

AccessDeniedException

The action cannot be performed because you do not have the required permission.

HTTP Status Code: 403

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

RenameTable

Service: Amazon S3 Tables

Renames a table or a namespace. For more information, see [S3 Tables](#) in the *Amazon Simple Storage Service User Guide*.

Permissions

You must have the `s3tables:RenameTable` permission to use this operation.

Request Syntax

```
PUT /tables/tableBucketARN/namespace/name/rename HTTP/1.1
Content-type: application/json
```

```
{
  "newName": "string",
  "newNamespaceName": "string",
  "versionToken": "string"
}
```

URI Request Parameters

The request uses the following URI parameters.

[name](#)

The current name of the table.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

[namespace](#)

The namespace associated with the table.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket.

Pattern: (arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})

Required: Yes

Request Body

The request accepts the following data in JSON format.

newName

The new name for the table.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: [0-9a-z_]*

Required: No

newNamespaceName

The new name for the namespace.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: [0-9a-z_]*

Required: No

versionToken

The version token of the table.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: No

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

TagResource

Service: Amazon S3 Tables

Applies one or more user-defined tags to an Amazon S3 Tables resource or updates existing tags. Each tag is a label consisting of a key and value pair. Tags can help you organize, track costs for, and control access to your resources. You can add up to 50 tags for each S3 resource.

Note

For a list of S3 resources that support tagging, see [Managing tags for Amazon S3 resources](#).

Permissions

For tables and table buckets, you must have the `s3tables:TagResource` permission to use this operation.

Request Syntax

```
POST /tag/resourceArn HTTP/1.1
Content-type: application/json
```

```
{
  "tags": {
    "string" : "string"
  }
}
```

URI Request Parameters

The request uses the following URI parameters.

resourceArn

The Amazon Resource Name (ARN) of the Amazon S3 Tables resource that you're applying tags to. The tagged resource can be a table bucket or a table. For a list of all S3 resources that support tagging, see [Managing tags for Amazon S3 resources](#).

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/.+`

Required: Yes

Request Body

The request accepts the following data in JSON format.

tags

The user-defined tag that you want to add to the specified S3 Tables resource. For more information, see [Tagging for cost allocation or attribute-based access control \(ABAC\)](#).

Type: String to string map

Key Length Constraints: Minimum length of 1. Maximum length of 128.

Value Length Constraints: Minimum length of 0. Maximum length of 256.

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

UntagResource

Service: Amazon S3 Tables

Removes the specified user-defined tags from an Amazon S3 Tables resource. You can pass one or more tag keys.

Note

For a list of S3 resources that support tagging, see [Managing tags for Amazon S3 resources](#).

Permissions

For tables and table buckets, you must have the `s3tables:UntagResource` permission to use this operation.

Request Syntax

```
DELETE /tag/resourceArn?tagKeys=tagKeys HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

resourceArn

The Amazon Resource Name (ARN) of the Amazon S3 Tables resource that you're removing tags from. The tagged resource can be a table bucket or a table. For a list of all S3 resources that support tagging, see [Managing tags for Amazon S3 resources](#).

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/.+`

Required: Yes

tagKeys

The array of tag keys that you're removing from the S3 Tables resource. For more information, see [Tagging for cost allocation or attribute-based access control \(ABAC\)](#).

Length Constraints: Minimum length of 1. Maximum length of 128.

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 204
```

Response Elements

If the action is successful, the service sends back an HTTP 204 response with an empty HTTP body.

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

UpdateTableMetadataLocation

Service: Amazon S3 Tables

Updates the metadata location for a table. The metadata location of a table must be an S3 URI that begins with the table's warehouse location. The metadata location for an Apache Iceberg table must end with `.metadata.json`, or if the metadata file is Gzip-compressed, `.metadata.json.gz`.

Permissions

You must have the `s3tables:UpdateTableMetadataLocation` permission to use this operation.

Request Syntax

```
PUT /tables/tableBucketARN/namespace/name/metadata-location HTTP/1.1
Content-type: application/json

{
  "metadataLocation": "string",
  "versionToken": "string"
}
```

URI Request Parameters

The request uses the following URI parameters.

name

The name of the table.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

namespace

The namespace of the table.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

tableBucketARN

The Amazon Resource Name (ARN) of the table bucket.

Pattern: `(arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})`

Required: Yes

Request Body

The request accepts the following data in JSON format.

metadataLocation

The new metadata location for the table.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: Yes

versionToken

The version token of the table.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: Yes

Response Syntax

```
HTTP/1.1 200
Content-type: application/json
```

```
{
  "metadataLocation": "string",
  "name": "string",
  "namespace": [ "string" ],
  "tableARN": "string",
  "versionToken": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

metadataLocation

The metadata location of the table.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

name

The name of the table.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: [0-9a-z_]*

namespace

The namespace the table is associated with.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: [0-9a-z_]*

tableARN

The Amazon Resource Name (ARN) of the table.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: (arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63}/table/[a-zA-Z0-9_-]{1,255})

versionToken

The version token of the table.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Errors

BadRequestException

The request is invalid or malformed.

HTTP Status Code: 400

ConflictException

The request failed because there is a conflict with a previous write. You can retry the request.

HTTP Status Code: 409

ForbiddenException

The caller isn't authorized to make the request.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource could not be found.

HTTP Status Code: 404

TooManyRequestsException

The limit on the number of requests per second was exceeded.

HTTP Status Code: 429

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

Amazon S3 Vectors

The following actions are supported by Amazon S3 Vectors:

- [CreateIndex](#)
- [CreateVectorBucket](#)
- [DeleteIndex](#)
- [DeleteVectorBucket](#)
- [DeleteVectorBucketPolicy](#)
- [DeleteVectors](#)
- [GetIndex](#)
- [GetVectorBucket](#)

- [GetVectorBucketPolicy](#)
- [GetVectors](#)
- [ListIndexes](#)
- [ListTagsForResource](#)
- [ListVectorBuckets](#)
- [ListVectors](#)
- [PutVectorBucketPolicy](#)
- [PutVectors](#)
- [QueryVectors](#)
- [TagResource](#)
- [UntagResource](#)

CreateIndex

Service: Amazon S3 Vectors

Creates a vector index within a vector bucket. To specify the vector bucket, you must use either the vector bucket name or the vector bucket Amazon Resource Name (ARN).

Permissions

You must have the `s3vectors:CreateIndex` permission to use this operation.

You must have the `s3vectors:TagResource` permission in addition to `s3vectors:CreateIndex` permission to create a vector index with tags.

Request Syntax

```
POST /CreateIndex HTTP/1.1
Content-type: application/json

{
  "dataType": "string",
  "dimension": number,
  "distanceMetric": "string",
  "encryptionConfiguration": {
    "kmsKeyArn": "string",
    "sseType": "string"
  },
  "indexName": "string",
  "metadataConfiguration": {
    "nonFilterableMetadataKeys": [ "string" ]
  },
  "tags": {
    "string" : "string"
  },
  "vectorBucketArn": "string",
  "vectorBucketName": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

dataType

The data type of the vectors to be inserted into the vector index.

Type: String

Valid Values: float32

Required: Yes

dimension

The dimensions of the vectors to be inserted into the vector index.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 4096.

Required: Yes

distanceMetric

The distance metric to be used for similarity search.

Type: String

Valid Values: euclidean | cosine

Required: Yes

encryptionConfiguration

The encryption configuration for a vector index. By default, if you don't specify, all new vectors in the vector index will use the encryption configuration of the vector bucket.

Type: [EncryptionConfiguration](#) object

Required: No

indexName

The name of the vector index to create.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: Yes

metadataConfiguration

The metadata configuration for the vector index.

Type: [MetadataConfiguration](#) object

Required: No

tags

An array of user-defined tags that you would like to apply to the vector index that you are creating. A tag is a key-value pair that you apply to your resources. Tags can help you organize, track costs, and control access to resources. For more information, see [Tagging for cost allocation or attribute-based access control \(ABAC\)](#).

Note

You must have the `s3vectors:TagResource` permission in addition to `s3vectors:CreateIndex` permission to create a vector index with tags.

Type: String to string map

Key Length Constraints: Minimum length of 1. Maximum length of 128.

Key Pattern: ([\p{L}\p{Z}\p{N}_ . : / = + \ - @] *)

Value Length Constraints: Minimum length of 0. Maximum length of 256.

Value Pattern: ([\p{L}\p{Z}\p{N}_ . : / = + \ - @] *)

Required: No

vectorBucketArn

The Amazon Resource Name (ARN) of the vector bucket to create the vector index in.

Type: String

Pattern: `arn:aws[-a-z0-9]*:s3vectors:[a-z0-9-]+:[0-9]{12}:bucket/[a-z0-9][a-z0-9-.]{1,61}[a-z0-9]`

Required: No

vectorBucketName

The name of the vector bucket to create the vector index in.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "indexArn": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

indexArn

The Amazon Resource Name (ARN) of the newly created vector index.

Type: String

Pattern: `arn:aws[-a-z0-9]*:s3vectors:[a-z0-9-]+:[0-9]{12}:bucket/[a-z0-9][a-z0-9-.]{1,61}[a-z0-9]/index/[a-z0-9][a-z0-9-.]{1,61}[a-z0-9]`

Errors

AccessDeniedException

Access denied.

HTTP Status Code: 403

ConflictException

The request failed because a vector bucket name or a vector index name already exists. Vector bucket names must be unique within your AWS account for each AWS Region. Vector index names must be unique within your vector bucket. Choose a different vector bucket name or vector index name, and try again.

HTTP Status Code: 409

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource can't be found.

HTTP Status Code: 404

RequestTimeoutException

The request timed out. Retry your request.

HTTP Status Code: 408

ServiceQuotaExceededException

Your request exceeds a service quota.

HTTP Status Code: 402

ServiceUnavailableException

The service is unavailable. Wait briefly and retry your request. If it continues to fail, increase your waiting time between retries.

HTTP Status Code: 503

TooManyRequestsException

The request was denied due to request throttling.

HTTP Status Code: 429

ValidationException

The requested action isn't valid.

fieldList

A list of specific validation failures that are encountered during input processing. Each entry in the list contains a path to the field that failed validation and a detailed message that explains why the validation failed. To satisfy multiple constraints, a field can appear multiple times in this list if it failed. You can use the information to identify and fix the specific validation issues in your request.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

CreateVectorBucket

Service: Amazon S3 Vectors

Creates a vector bucket in the AWS Region that you want your bucket to be in.

Permissions

You must have the `s3vectors:CreateVectorBucket` permission to use this operation.

You must have the `s3vectors:TagResource` permission in addition to `s3vectors:CreateVectorBucket` permission to create a vector bucket with tags.

Request Syntax

```
POST /CreateVectorBucket HTTP/1.1
```

```
Content-type: application/json
```

```
{
  "encryptionConfiguration": {
    "kmsKeyArn": "string",
    "sseType": "string"
  },
  "tags": {
    "string" : "string"
  },
  "vectorBucketName": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

encryptionConfiguration

The encryption configuration for the vector bucket. By default, if you don't specify, all new vectors in Amazon S3 vector buckets use server-side encryption with Amazon S3 managed keys (SSE-S3), specifically AES256.

Type: [EncryptionConfiguration](#) object

Required: No

[tags](#)

An array of user-defined tags that you would like to apply to the vector bucket that you are creating. A tag is a key-value pair that you apply to your resources. Tags can help you organize and control access to resources. For more information, see [Tagging for cost allocation or attribute-based access control \(ABAC\)](#).

Note

You must have the `s3vectors:TagResource` permission in addition to `s3vectors:CreateVectorBucket` permission to create a vector bucket with tags.

Type: String to string map

Key Length Constraints: Minimum length of 1. Maximum length of 128.

Key Pattern: ([\p{L}\p{Z}\p{N}_ : / = + \ - @] *)

Value Length Constraints: Minimum length of 0. Maximum length of 256.

Value Pattern: ([\p{L}\p{Z}\p{N}_ : / = + \ - @] *)

Required: No

[vectorBucketName](#)

The name of the vector bucket to create.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

```
Content-type: application/json

{
  "vectorBucketArn": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[vectorBucketArn](#)

The Amazon Resource Name (ARN) of the newly created vector bucket.

Type: String

Pattern: `arn:aws[-a-z0-9]*:s3vectors:[a-z0-9-]+:[0-9]{12}:bucket/[a-z0-9][a-z0-9- .]{1,61}[a-z0-9]`

Errors

AccessDeniedException

Access denied.

HTTP Status Code: 403

ConflictException

The request failed because a vector bucket name or a vector index name already exists. Vector bucket names must be unique within your AWS account for each AWS Region. Vector index names must be unique within your vector bucket. Choose a different vector bucket name or vector index name, and try again.

HTTP Status Code: 409

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

RequestTimeoutException

The request timed out. Retry your request.

HTTP Status Code: 408

ServiceQuotaExceededException

Your request exceeds a service quota.

HTTP Status Code: 402

ServiceUnavailableException

The service is unavailable. Wait briefly and retry your request. If it continues to fail, increase your waiting time between retries.

HTTP Status Code: 503

TooManyRequestsException

The request was denied due to request throttling.

HTTP Status Code: 429

ValidationException

The requested action isn't valid.

fieldList

A list of specific validation failures that are encountered during input processing. Each entry in the list contains a path to the field that failed validation and a detailed message that explains why the validation failed. To satisfy multiple constraints, a field can appear multiple times in this list if it failed. You can use the information to identify and fix the specific validation issues in your request.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)

- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteIndex

Service: Amazon S3 Vectors

Deletes a vector index. To specify the vector index, you can either use both the vector bucket name and vector index name, or use the vector index Amazon Resource Name (ARN).

Permissions

You must have the `s3vectors:DeleteIndex` permission to use this operation.

Request Syntax

```
POST /DeleteIndex HTTP/1.1
Content-type: application/json

{
  "indexArn": "string",
  "indexName": "string",
  "vectorBucketName": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

indexArn

The ARN of the vector index to delete.

Type: String

Pattern: `arn:aws[-a-z0-9]*:s3vectors:[a-z0-9-]+:[0-9]{12}:bucket/[a-z0-9][a-z0-9-]{1,61}[a-z0-9]/index/[a-z0-9][a-z0-9-]{1,61}[a-z0-9]`

Required: No

indexName

The name of the vector index to delete.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: No

vectorBucketName

The name of the vector bucket that contains the vector index.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: No

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

AccessDeniedException

Access denied.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource can't be found.

HTTP Status Code: 404

RequestTimeoutException

The request timed out. Retry your request.

HTTP Status Code: 408

ServiceUnavailableException

The service is unavailable. Wait briefly and retry your request. If it continues to fail, increase your waiting time between retries.

HTTP Status Code: 503

TooManyRequestsException

The request was denied due to request throttling.

HTTP Status Code: 429

ValidationException

The requested action isn't valid.

fieldList

A list of specific validation failures that are encountered during input processing. Each entry in the list contains a path to the field that failed validation and a detailed message that explains why the validation failed. To satisfy multiple constraints, a field can appear multiple times in this list if it failed. You can use the information to identify and fix the specific validation issues in your request.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteVectorBucket

Service: Amazon S3 Vectors

Deletes a vector bucket. All vector indexes in the vector bucket must be deleted before the vector bucket can be deleted. To perform this operation, you must use either the vector bucket name or the vector bucket Amazon Resource Name (ARN).

Permissions

You must have the `s3vectors:DeleteVectorBucket` permission to use this operation.

Request Syntax

```
POST /DeleteVectorBucket HTTP/1.1
Content-type: application/json

{
  "vectorBucketArn": "string",
  "vectorBucketName": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

vectorBucketArn

The ARN of the vector bucket to delete.

Type: String

Pattern: `arn:aws[-a-z0-9]*:s3vectors:[a-z0-9-]+:[0-9]{12}:bucket/[a-z0-9][a-z0-9-.]{1,61}[a-z0-9]`

Required: No

vectorBucketName

The name of the vector bucket to delete.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: No

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

AccessDeniedException

Access denied.

HTTP Status Code: 403

ConflictException

The request failed because a vector bucket name or a vector index name already exists. Vector bucket names must be unique within your AWS account for each AWS Region. Vector index names must be unique within your vector bucket. Choose a different vector bucket name or vector index name, and try again.

HTTP Status Code: 409

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource can't be found.

HTTP Status Code: 404

RequestTimeoutException

The request timed out. Retry your request.

HTTP Status Code: 408

ServiceUnavailableException

The service is unavailable. Wait briefly and retry your request. If it continues to fail, increase your waiting time between retries.

HTTP Status Code: 503

TooManyRequestsException

The request was denied due to request throttling.

HTTP Status Code: 429

ValidationException

The requested action isn't valid.

fieldList

A list of specific validation failures that are encountered during input processing. Each entry in the list contains a path to the field that failed validation and a detailed message that explains why the validation failed. To satisfy multiple constraints, a field can appear multiple times in this list if it failed. You can use the information to identify and fix the specific validation issues in your request.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteVectorBucketPolicy

Service: Amazon S3 Vectors

Deletes a vector bucket policy. To specify the bucket, you must use either the vector bucket name or the vector bucket Amazon Resource Name (ARN).

Permissions

You must have the `s3vectors:DeleteVectorBucketPolicy` permission to use this operation.

Request Syntax

```
POST /DeleteVectorBucketPolicy HTTP/1.1
Content-type: application/json
```

```
{
  "vectorBucketArn": "string",
  "vectorBucketName": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

[vectorBucketArn](#)

The ARN of the vector bucket to delete the policy from.

Type: String

Pattern: `arn:aws[-a-z0-9]*:s3vectors:[a-z0-9-]+:[0-9]{12}:bucket/[a-z0-9][a-z0-9-.]{1,61}[a-z0-9]`

Required: No

[vectorBucketName](#)

The name of the vector bucket to delete the policy from.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: No

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

AccessDeniedException

Access denied.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource can't be found.

HTTP Status Code: 404

RequestTimeoutException

The request timed out. Retry your request.

HTTP Status Code: 408

ServiceUnavailableException

The service is unavailable. Wait briefly and retry your request. If it continues to fail, increase your waiting time between retries.

HTTP Status Code: 503

TooManyRequestsException

The request was denied due to request throttling.

HTTP Status Code: 429

ValidationException

The requested action isn't valid.

fieldList

A list of specific validation failures that are encountered during input processing. Each entry in the list contains a path to the field that failed validation and a detailed message that explains why the validation failed. To satisfy multiple constraints, a field can appear multiple times in this list if it failed. You can use the information to identify and fix the specific validation issues in your request.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteVectors

Service: Amazon S3 Vectors

Deletes one or more vectors in a vector index. To specify the vector index, you can either use both the vector bucket name and vector index name, or use the vector index Amazon Resource Name (ARN).

Permissions

You must have the `s3vectors:DeleteVectors` permission to use this operation.

Request Syntax

```
POST /DeleteVectors HTTP/1.1
Content-type: application/json

{
  "indexArn": "string",
  "indexName": "string",
  "keys": [ "string" ],
  "vectorBucketName": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

[indexArn](#)

The ARN of the vector index that contains a vector you want to delete.

Type: String

Pattern: `arn:aws[-a-z0-9]*:s3vectors:[a-z0-9-]+:[0-9]{12}:bucket/[a-z0-9][a-z0-9-.]{1,61}[a-z0-9]/index/[a-z0-9][a-z0-9-.]{1,61}[a-z0-9]`

Required: No

indexName

The name of the vector index that contains a vector you want to delete.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: No

keys

The keys of the vectors to delete.

Type: Array of strings

Array Members: Minimum number of 1 item. Maximum number of 500 items.

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

vectorBucketName

The name of the vector bucket that contains the vector index.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: No

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

AccessDeniedException

Access denied.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

KmsDisabledException

The specified AWS KMS key isn't enabled.

HTTP Status Code: 400

KmsInvalidKeyUsageException

The request was rejected for one of the following reasons:

- The KeyUsage value of the KMS key is incompatible with the API operation.
- The encryption algorithm or signing algorithm specified for the operation is incompatible with the type of key material in the KMS key (KeySpec).

For more information, see [InvalidKeyUsageException](#) in the *AWS Key Management Service API Reference*.

HTTP Status Code: 400

KmsInvalidStateException

The key state of the KMS key isn't compatible with the operation.

For more information, see [KMSInvalidStateException](#) in the *AWS Key Management Service API Reference*.

HTTP Status Code: 400

KmsNotFoundException

The KMS key can't be found.

HTTP Status Code: 400

NotFoundException

The request was rejected because the specified resource can't be found.

HTTP Status Code: 404

RequestTimeoutException

The request timed out. Retry your request.

HTTP Status Code: 408

ServiceUnavailableException

The service is unavailable. Wait briefly and retry your request. If it continues to fail, increase your waiting time between retries.

HTTP Status Code: 503

TooManyRequestsException

The request was denied due to request throttling.

HTTP Status Code: 429

ValidationException

The requested action isn't valid.

fieldList

A list of specific validation failures that are encountered during input processing. Each entry in the list contains a path to the field that failed validation and a detailed message that explains why the validation failed. To satisfy multiple constraints, a field can appear multiple times in this list if it failed. You can use the information to identify and fix the specific validation issues in your request.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetIndex

Service: Amazon S3 Vectors

Returns vector index attributes. To specify the vector index, you can either use both the vector bucket name and the vector index name, or use the vector index Amazon Resource Name (ARN).

Permissions

You must have the `s3vectors:GetIndex` permission to use this operation.

Request Syntax

```
POST /GetIndex HTTP/1.1
Content-type: application/json

{
  "indexArn": "string",
  "indexName": "string",
  "vectorBucketName": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

indexArn

The ARN of the vector index.

Type: String

Pattern: `arn:aws[-a-z0-9]*:s3vectors:[a-z0-9-]+:[0-9]{12}:bucket/[a-z0-9][a-z0-9-]{1,61}[a-z0-9]/index/[a-z0-9][a-z0-9-]{1,61}[a-z0-9]`

Required: No

indexName

The name of the vector index.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: No

vectorBucketName

The name of the vector bucket that contains the vector index.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "index": {
    "creationTime": number,
    "dataType": "string",
    "dimension": number,
    "distanceMetric": "string",
    "encryptionConfiguration": {
      "kmsKeyArn": "string",
      "sseType": "string"
    },
    "indexArn": "string",
    "indexName": "string",
    "metadataConfiguration": {
      "nonFilterableMetadataKeys": [ "string" ]
    },
    "vectorBucketName": "string"
  }
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[index](#)

The attributes of the vector index.

Type: [Index](#) object

Errors

AccessDeniedException

Access denied.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource can't be found.

HTTP Status Code: 404

RequestTimeoutException

The request timed out. Retry your request.

HTTP Status Code: 408

ServiceUnavailableException

The service is unavailable. Wait briefly and retry your request. If it continues to fail, increase your waiting time between retries.

HTTP Status Code: 503

TooManyRequestsException

The request was denied due to request throttling.

HTTP Status Code: 429

ValidationException

The requested action isn't valid.

fieldList

A list of specific validation failures that are encountered during input processing. Each entry in the list contains a path to the field that failed validation and a detailed message that explains why the validation failed. To satisfy multiple constraints, a field can appear multiple times in this list if it failed. You can use the information to identify and fix the specific validation issues in your request.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetVectorBucket

Service: Amazon S3 Vectors

Returns vector bucket attributes. To specify the bucket, you must use either the vector bucket name or the vector bucket Amazon Resource Name (ARN).

Permissions

You must have the `s3vectors:GetVectorBucket` permission to use this operation.

Request Syntax

```
POST /GetVectorBucket HTTP/1.1
Content-type: application/json

{
  "vectorBucketArn": "string",
  "vectorBucketName": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

vectorBucketArn

The ARN of the vector bucket to retrieve information about.

Type: String

Pattern: `arn:aws[-a-z0-9]*:s3vectors:[a-z0-9-]+:[0-9]{12}:bucket/[a-z0-9][a-z0-9-]{1,61}[a-z0-9]`

Required: No

vectorBucketName

The name of the vector bucket to retrieve information about.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "vectorBucket": {
    "creationTime": number,
    "encryptionConfiguration": {
      "kmsKeyArn": "string",
      "sseType": "string"
    },
    "vectorBucketArn": "string",
    "vectorBucketName": "string"
  }
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

vectorBucket

The attributes of the vector bucket.

Type: [VectorBucket](#) object

Errors

AccessDeniedException

Access denied.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource can't be found.

HTTP Status Code: 404

RequestTimeoutException

The request timed out. Retry your request.

HTTP Status Code: 408

ServiceUnavailableException

The service is unavailable. Wait briefly and retry your request. If it continues to fail, increase your waiting time between retries.

HTTP Status Code: 503

TooManyRequestsException

The request was denied due to request throttling.

HTTP Status Code: 429

ValidationException

The requested action isn't valid.

fieldList

A list of specific validation failures that are encountered during input processing. Each entry in the list contains a path to the field that failed validation and a detailed message that explains why the validation failed. To satisfy multiple constraints, a field can appear multiple times in this list if it failed. You can use the information to identify and fix the specific validation issues in your request.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetVectorBucketPolicy

Service: Amazon S3 Vectors

Gets details about a vector bucket policy. To specify the bucket, you must use either the vector bucket name or the vector bucket Amazon Resource Name (ARN).

Permissions

You must have the `s3vectors:GetVectorBucketPolicy` permission to use this operation.

Request Syntax

```
POST /GetVectorBucketPolicy HTTP/1.1
Content-type: application/json

{
  "vectorBucketArn": "string",
  "vectorBucketName": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

vectorBucketArn

The ARN of the vector bucket.

Type: String

Pattern: `arn:aws[-a-z0-9]*:s3vectors:[a-z0-9-]+:[0-9]{12}:bucket/[a-z0-9][a-z0-9-.]{1,61}[a-z0-9]`

Required: No

vectorBucketName

The name of the vector bucket.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "policy": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[policy](#)

The JSON that defines the policy.

Type: String

Errors

AccessDeniedException

Access denied.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource can't be found.

HTTP Status Code: 404

RequestTimeoutException

The request timed out. Retry your request.

HTTP Status Code: 408

ServiceUnavailableException

The service is unavailable. Wait briefly and retry your request. If it continues to fail, increase your waiting time between retries.

HTTP Status Code: 503

TooManyRequestsException

The request was denied due to request throttling.

HTTP Status Code: 429

ValidationException

The requested action isn't valid.

fieldList

A list of specific validation failures that are encountered during input processing. Each entry in the list contains a path to the field that failed validation and a detailed message that explains why the validation failed. To satisfy multiple constraints, a field can appear multiple times in this list if it failed. You can use the information to identify and fix the specific validation issues in your request.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)

- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetVectors

Service: Amazon S3 Vectors

Returns vector attributes. To specify the vector index, you can either use both the vector bucket name and the vector index name, or use the vector index Amazon Resource Name (ARN).

Permissions

You must have the `s3vectors:GetVectors` permission to use this operation.

Request Syntax

```
POST /GetVectors HTTP/1.1
Content-type: application/json

{
  "indexArn": "string",
  "indexName": "string",
  "keys": [ "string" ],
  "returnData": boolean,
  "returnMetadata": boolean,
  "vectorBucketName": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

indexArn

The ARN of the vector index.

Type: String

Pattern: `arn:aws[-a-z0-9]*:s3vectors:[a-z0-9-]+:[0-9]{12}:bucket/[a-z0-9][a-z0-9- .]{1,61}[a-z0-9]/index/[a-z0-9][a-z0-9- .]{1,61}[a-z0-9]`

Required: No

indexName

The name of the vector index.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: No

keys

The names of the vectors you want to return attributes for.

Type: Array of strings

Array Members: Minimum number of 1 item. Maximum number of 100 items.

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

returnData

Indicates whether to include the vector data in the response. The default value is `false`.

Type: Boolean

Required: No

returnMetadata

Indicates whether to include metadata in the response. The default value is `false`.

Type: Boolean

Required: No

vectorBucketName

The name of the vector bucket that contains the vector index.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "vectors": [
    {
      "data": { ... },
      "key": "string",
      "metadata": JSON value
    }
  ]
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

vectors

The attributes of the vectors.

Type: Array of [GetOutputVector](#) objects

Errors

AccessDeniedException

Access denied.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

KmsDisabledException

The specified AWS KMS key isn't enabled.

HTTP Status Code: 400

KmsInvalidKeyUsageException

The request was rejected for one of the following reasons:

- The KeyUsage value of the KMS key is incompatible with the API operation.
- The encryption algorithm or signing algorithm specified for the operation is incompatible with the type of key material in the KMS key (KeySpec).

For more information, see [InvalidKeyUsageException](#) in the *AWS Key Management Service API Reference*.

HTTP Status Code: 400

KmsInvalidStateException

The key state of the KMS key isn't compatible with the operation.

For more information, see [KMSInvalidStateException](#) in the *AWS Key Management Service API Reference*.

HTTP Status Code: 400

KmsNotFoundException

The KMS key can't be found.

HTTP Status Code: 400

NotFoundException

The request was rejected because the specified resource can't be found.

HTTP Status Code: 404

RequestTimeoutException

The request timed out. Retry your request.

HTTP Status Code: 408

ServiceUnavailableException

The service is unavailable. Wait briefly and retry your request. If it continues to fail, increase your waiting time between retries.

HTTP Status Code: 503

TooManyRequestsException

The request was denied due to request throttling.

HTTP Status Code: 429

ValidationException

The requested action isn't valid.

fieldList

A list of specific validation failures that are encountered during input processing. Each entry in the list contains a path to the field that failed validation and a detailed message that explains why the validation failed. To satisfy multiple constraints, a field can appear multiple times in this list if it failed. You can use the information to identify and fix the specific validation issues in your request.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)

- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListIndexes

Service: Amazon S3 Vectors

Returns a list of all the vector indexes within the specified vector bucket. To specify the bucket, you must use either the vector bucket name or the vector bucket Amazon Resource Name (ARN).

Permissions

You must have the `s3vectors:ListIndexes` permission to use this operation.

Request Syntax

```
POST /ListIndexes HTTP/1.1
Content-type: application/json

{
  "maxResults": number,
  "nextToken": "string",
  "prefix": "string",
  "vectorBucketArn": "string",
  "vectorBucketName": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

maxResults

The maximum number of items to be returned in the response.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 500.

Required: No

nextToken

The previous pagination token.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Required: No

prefix

Limits the response to vector indexes that begin with the specified prefix.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 63.

Required: No

vectorBucketArn

The ARN of the vector bucket that contains the vector indexes.

Type: String

Pattern: `arn:aws[-a-z0-9]*:s3vectors:[a-z0-9-]+:[0-9]{12}:bucket/[a-z0-9][a-z0-9- .]{1,61}[a-z0-9]`

Required: No

vectorBucketName

The name of the vector bucket that contains the vector indexes.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json
```

```
{
  "indexes": [
    {
      "creationTime": number,
      "indexArn": "string",
      "indexName": "string",
      "vectorBucketName": "string"
    }
  ],
  "nextToken": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

indexes

The attributes of the vector indexes

Type: Array of [IndexSummary](#) objects

nextToken

The next pagination token.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Errors

AccessDeniedException

Access denied.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource can't be found.

HTTP Status Code: 404

RequestTimeoutException

The request timed out. Retry your request.

HTTP Status Code: 408

ServiceUnavailableException

The service is unavailable. Wait briefly and retry your request. If it continues to fail, increase your waiting time between retries.

HTTP Status Code: 503

TooManyRequestsException

The request was denied due to request throttling.

HTTP Status Code: 429

ValidationException

The requested action isn't valid.

fieldList

A list of specific validation failures that are encountered during input processing. Each entry in the list contains a path to the field that failed validation and a detailed message that explains why the validation failed. To satisfy multiple constraints, a field can appear multiple times in this list if it failed. You can use the information to identify and fix the specific validation issues in your request.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListTagsForResource

Service: Amazon S3 Vectors

Lists all of the tags applied to a specified Amazon S3 Vectors resource. Each tag is a label consisting of a key and value pair. Tags can help you organize, track costs for, and control access to resources.

Note

For a list of S3 resources that support tagging, see [Managing tags for Amazon S3 resources](#).

Permissions

For vector buckets and vector indexes, you must have the `s3vectors:ListTagsForResource` permission to use this operation.

Request Syntax

```
GET /tags/resourceArn HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

resourceArn

The Amazon Resource Name (ARN) of the Amazon S3 Vectors resource that you want to list tags for. The tagged resource can be a vector bucket or a vector index.

Length Constraints: Minimum length of 0. Maximum length of 1011.

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

```
Content-type: application/json

{
  "tags": {
    "string" : "string"
  }
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

tags

The user-defined tags that are applied to the S3 Vectors resource. For more information, see [Tagging for cost allocation or attribute-based access control \(ABAC\)](#).

Type: String to string map

Key Length Constraints: Minimum length of 1. Maximum length of 128.

Key Pattern: ([\p{L}\p{Z}\p{N}_. :/=+\-@] *)

Value Length Constraints: Minimum length of 0. Maximum length of 256.

Value Pattern: ([\p{L}\p{Z}\p{N}_. :/=+\-@] *)

Errors

AccessDeniedException

Access denied.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource can't be found.

HTTP Status Code: 404

RequestTimeoutException

The request timed out. Retry your request.

HTTP Status Code: 408

ServiceUnavailableException

The service is unavailable. Wait briefly and retry your request. If it continues to fail, increase your waiting time between retries.

HTTP Status Code: 503

TooManyRequestsException

The request was denied due to request throttling.

HTTP Status Code: 429

ValidationException

The requested action isn't valid.

fieldList

A list of specific validation failures that are encountered during input processing. Each entry in the list contains a path to the field that failed validation and a detailed message that explains why the validation failed. To satisfy multiple constraints, a field can appear multiple times in this list if it failed. You can use the information to identify and fix the specific validation issues in your request.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)

- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListVectorBuckets

Service: Amazon S3 Vectors

Returns a list of all the vector buckets that are owned by the authenticated sender of the request.

Permissions

You must have the `s3vectors:ListVectorBuckets` permission to use this operation.

Request Syntax

```
POST /ListVectorBuckets HTTP/1.1
Content-type: application/json
```

```
{
  "maxResults": number,
  "nextToken": "string",
  "prefix": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

maxResults

The maximum number of vector buckets to be returned in the response.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 500.

Required: No

nextToken

The previous pagination token.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Required: No

prefix

Limits the response to vector buckets that begin with the specified prefix.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 63.

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "nextToken": "string",
  "vectorBuckets": [
    {
      "creationTime": number,
      "vectorBucketArn": "string",
      "vectorBucketName": "string"
    }
  ]
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

nextToken

The element is included in the response when there are more buckets to be listed with pagination.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

[vectorBuckets](#)

The list of vector buckets owned by the requester.

Type: Array of [VectorBucketSummary](#) objects

Errors

AccessDeniedException

Access denied.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

RequestTimeoutException

The request timed out. Retry your request.

HTTP Status Code: 408

ServiceUnavailableException

The service is unavailable. Wait briefly and retry your request. If it continues to fail, increase your waiting time between retries.

HTTP Status Code: 503

TooManyRequestsException

The request was denied due to request throttling.

HTTP Status Code: 429

ValidationException

The requested action isn't valid.

fieldList

A list of specific validation failures that are encountered during input processing. Each entry in the list contains a path to the field that failed validation and a detailed message that explains why the validation failed. To satisfy multiple constraints, a field can appear multiple times in this list if it failed. You can use the information to identify and fix the specific validation issues in your request.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListVectors

Service: Amazon S3 Vectors

List vectors in the specified vector index. To specify the vector index, you can either use both the vector bucket name and the vector index name, or use the vector index Amazon Resource Name (ARN).

ListVectors operations proceed sequentially; however, for faster performance on a large number of vectors in a vector index, applications can request a parallel ListVectors operation by providing the segmentCount and segmentIndex parameters.

Permissions

You must have the `s3vectors:ListVectors` permission to use this operation. Additional permissions are required based on the request parameters you specify:

- With only `s3vectors:ListVectors` permission, you can list vector keys when `returnData` and `returnMetadata` are both set to `false` or not specified..
- If you set `returnData` or `returnMetadata` to `true`, you must have both `s3vectors:ListVectors` and `s3vectors:GetVectors` permissions. The request fails with a 403 Forbidden error if you request vector data or metadata without the `s3vectors:GetVectors` permission.

Request Syntax

```
POST /ListVectors HTTP/1.1
Content-type: application/json
```

```
{
  "indexArn": "string",
  "indexName": "string",
  "maxResults": number,
  "nextToken": "string",
  "returnData": boolean,
  "returnMetadata": boolean,
  "segmentCount": number,
  "segmentIndex": number,
  "vectorBucketName": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

indexArn

The Amazon resource Name (ARN) of the vector index.

Type: String

Pattern: `arn:aws[-a-z0-9]*:s3vectors:[a-z0-9-]+:[0-9]{12}:bucket/[a-z0-9][a-z0-9-.]{1,61}[a-z0-9]/index/[a-z0-9][a-z0-9-.]{1,61}[a-z0-9]`

Required: No

indexName

The name of the vector index.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: No

maxResults

The maximum number of vectors to return on a page.

If you don't specify `maxResults`, the `ListVectors` operation uses a default value of 500.

If the processed dataset size exceeds 1 MB before reaching the `maxResults` value, the operation stops and returns the vectors that are retrieved up to that point, along with a `nextToken` that you can use in a subsequent request to retrieve the next set of results.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 1000.

Required: No

nextToken

Pagination token from a previous request. The value of this field is empty for an initial request.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: No

returnData

If true, the vector data of each vector will be included in the response. The default value is false.

Type: Boolean

Required: No

returnMetadata

If true, the metadata associated with each vector will be included in the response. The default value is false.

Type: Boolean

Required: No

segmentCount

For a parallel `ListVectors` request, `segmentCount` represents the total number of vector segments into which the `ListVectors` operation will be divided. The value of `segmentCount` corresponds to the number of application workers that will perform the parallel `ListVectors` operation. For example, if you want to use four application threads to list vectors in a vector index, specify a `segmentCount` value of 4.

If you specify a `segmentCount` value of 1, the `ListVectors` operation will be sequential rather than parallel.

If you specify `segmentCount`, you must also specify `segmentIndex`.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 16.

Required: No

segmentIndex

For a parallel `ListVectors` request, `segmentIndex` is the index of the segment from which to list vectors in the current request. It identifies an individual segment to be listed by an application worker.

Segment IDs are zero-based, so the first segment is always 0. For example, if you want to use four application threads to list vectors in a vector index, then the first thread specifies a `segmentIndex` value of 0, the second thread specifies 1, and so on.

The value of `segmentIndex` must be less than the value provided for `segmentCount`.

If you provide `segmentIndex`, you must also provide `segmentCount`.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 15.

Required: No

vectorBucketName

The name of the vector bucket.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json
```

```
{
  "nextToken": "string",
  "vectors": [
    {
      "data": { ... },
      "key": "string",
```

```
    "metadata": JSON value
  }
]
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[nextToken](#)

Pagination token to be used in the subsequent request. The field is empty if no further pagination is required.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

[vectors](#)

Vectors in the current segment.

Type: Array of [ListOutputVector](#) objects

Errors

AccessDeniedException

Access denied.

HTTP Status Code: 403

InternalServerError

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource can't be found.

HTTP Status Code: 404

RequestTimeoutException

The request timed out. Retry your request.

HTTP Status Code: 408

ServiceUnavailableException

The service is unavailable. Wait briefly and retry your request. If it continues to fail, increase your waiting time between retries.

HTTP Status Code: 503

TooManyRequestsException

The request was denied due to request throttling.

HTTP Status Code: 429

ValidationException

The requested action isn't valid.

fieldList

A list of specific validation failures that are encountered during input processing. Each entry in the list contains a path to the field that failed validation and a detailed message that explains why the validation failed. To satisfy multiple constraints, a field can appear multiple times in this list if it failed. You can use the information to identify and fix the specific validation issues in your request.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutVectorBucketPolicy

Service: Amazon S3 Vectors

Creates a bucket policy for a vector bucket. To specify the bucket, you must use either the vector bucket name or the vector bucket Amazon Resource Name (ARN).

Permissions

You must have the `s3vectors:PutVectorBucketPolicy` permission to use this operation.

Request Syntax

```
POST /PutVectorBucketPolicy HTTP/1.1
Content-type: application/json
```

```
{
  "policy": "string",
  "vectorBucketArn": "string",
  "vectorBucketName": "string"
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

policy

The JSON that defines the policy. For more information about bucket policies for S3 Vectors, see [Managing vector bucket policies](#) in the *Amazon S3 User Guide*.

Type: String

Required: Yes

vectorBucketArn

The Amazon Resource Name (ARN) of the vector bucket.

Type: String

Pattern: `arn:aws[-a-z0-9]*:s3vectors:[a-z0-9-]+:[0-9]{12}:bucket/[a-z0-9][a-z0-9-.]{1,61}[a-z0-9]`

Required: No

vectorBucketName

The name of the vector bucket.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: No

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

AccessDeniedException

Access denied.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource can't be found.

HTTP Status Code: 404

RequestTimeoutException

The request timed out. Retry your request.

HTTP Status Code: 408

ServiceUnavailableException

The service is unavailable. Wait briefly and retry your request. If it continues to fail, increase your waiting time between retries.

HTTP Status Code: 503

TooManyRequestsException

The request was denied due to request throttling.

HTTP Status Code: 429

ValidationException

The requested action isn't valid.

fieldList

A list of specific validation failures that are encountered during input processing. Each entry in the list contains a path to the field that failed validation and a detailed message that explains why the validation failed. To satisfy multiple constraints, a field can appear multiple times in this list if it failed. You can use the information to identify and fix the specific validation issues in your request.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

PutVectors

Service: Amazon S3 Vectors

Adds one or more vectors to a vector index. To specify the vector index, you can either use both the vector bucket name and the vector index name, or use the vector index Amazon Resource Name (ARN).

For more information about limits, see [Limitations and restrictions](#) in the *Amazon S3 User Guide*.

Note

When inserting vector data into your vector index, you must provide the vector data as float32 (32-bit floating point) values. If you pass higher-precision values to an AWS SDK, S3 Vectors converts the values to 32-bit floating point before storing them, and GetVectors, ListVectors, and QueryVectors operations return the float32 values. Different AWS SDKs may have different default numeric types, so ensure your vectors are properly formatted as float32 values regardless of which SDK you're using. For example, in Python, use `numpy.float32` or explicitly cast your values.

Permissions

You must have the `s3vectors:PutVectors` permission to use this operation.

Request Syntax

```
POST /PutVectors HTTP/1.1
Content-type: application/json

{
  "indexArn": "string",
  "indexName": "string",
  "vectorBucketName": "string",
  "vectors": [
    {
      "data": { ... },
      "key": "string",
      "metadata": JSON value
    }
  ]
}
```

```
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

indexArn

The ARN of the vector index where you want to write vectors.

Type: String

Pattern: `arn:aws[-a-z0-9]*:s3vectors:[a-z0-9-]+:[0-9]{12}:bucket/[a-z0-9][a-z0-9-]{1,61}[a-z0-9]/index/[a-z0-9][a-z0-9-]{1,61}[a-z0-9]`

Required: No

indexName

The name of the vector index where you want to write vectors.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: No

vectorBucketName

The name of the vector bucket that contains the vector index.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: No

vectors

The vectors to add to a vector index. The number of vectors in a single request must not exceed the resource capacity, otherwise the request will be rejected with the error

ServiceUnavailableException with the error message "Currently unable to handle the request".

Type: Array of [PutInputVector](#) objects

Array Members: Minimum number of 1 item. Maximum number of 500 items.

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

AccessDeniedException

Access denied.

HTTP Status Code: 403

InternalServerError

The request failed due to an internal server error.

HTTP Status Code: 500

KmsDisabledException

The specified AWS KMS key isn't enabled.

HTTP Status Code: 400

KmsInvalidKeyUsageException

The request was rejected for one of the following reasons:

- The KeyUsage value of the KMS key is incompatible with the API operation.
- The encryption algorithm or signing algorithm specified for the operation is incompatible with the type of key material in the KMS key (KeySpec).

For more information, see [InvalidKeyUsageException](#) in the *AWS Key Management Service API Reference*.

HTTP Status Code: 400

KmsInvalidStateException

The key state of the KMS key isn't compatible with the operation.

For more information, see [KMSInvalidStateException](#) in the *AWS Key Management Service API Reference*.

HTTP Status Code: 400

KmsNotFoundException

The KMS key can't be found.

HTTP Status Code: 400

NotFoundException

The request was rejected because the specified resource can't be found.

HTTP Status Code: 404

RequestTimeoutException

The request timed out. Retry your request.

HTTP Status Code: 408

ServiceQuotaExceededException

Your request exceeds a service quota.

HTTP Status Code: 402

ServiceUnavailableException

The service is unavailable. Wait briefly and retry your request. If it continues to fail, increase your waiting time between retries.

HTTP Status Code: 503

TooManyRequestsException

The request was denied due to request throttling.

HTTP Status Code: 429

ValidationException

The requested action isn't valid.

fieldList

A list of specific validation failures that are encountered during input processing. Each entry in the list contains a path to the field that failed validation and a detailed message that explains why the validation failed. To satisfy multiple constraints, a field can appear multiple times in this list if it failed. You can use the information to identify and fix the specific validation issues in your request.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

QueryVectors

Service: Amazon S3 Vectors

Performs an approximate nearest neighbor search query in a vector index using a query vector. By default, it returns the keys of approximate nearest neighbors. You can optionally include the computed distance (between the query vector and each vector in the response) and metadata of each vector in the response.

To specify the vector index, you can either use both the vector bucket name and the vector index name, or use the vector index Amazon Resource Name (ARN).

Permissions

You must have the `s3vectors:QueryVectors` permission to use this operation. Additional permissions are required based on the request parameters you specify:

- With only `s3vectors:QueryVectors` permission, you can retrieve vector keys of approximate nearest neighbors and computed distances between these vectors. This permission is sufficient only when you don't set any metadata filters and don't request metadata (by keeping the `returnMetadata` parameter set to `false` or not specified).
- If you specify a metadata filter or set `returnMetadata` to `true`, you must have both `s3vectors:QueryVectors` and `s3vectors:GetVectors` permissions. The request fails with a `403 Forbidden` error if you request metadata filtering or metadata without the `s3vectors:GetVectors` permission.

Request Syntax

```
POST /QueryVectors HTTP/1.1
Content-type: application/json
```

```
{
  "filter": JSON value,
  "indexArn": "string",
  "indexName": "string",
  "nextToken": "string",
  "queryVector": { ... },
  "returnDistance": boolean,
  "returnMetadata": boolean,
  "topK": number,
  "vectorBucketName": "string"
```

```
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

filter

Metadata filter to apply during the query. For more information about metadata keys, see [Metadata filtering](#) in the *Amazon S3 User Guide*.

Type: JSON value

Required: No

indexArn

The ARN of the vector index that you want to query.

Type: String

Pattern: `arn:aws[-a-z0-9]*:s3vectors:[a-z0-9-]+:[0-9]{12}:bucket/[a-z0-9][a-z0-9-.]{1,61}[a-z0-9]/index/[a-z0-9][a-z0-9-.]{1,61}[a-z0-9]`

Required: No

indexName

The name of the vector index that you want to query.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: No

nextToken

Pagination token from a previous request. The value of this field is empty for an initial request.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 4096.

Required: No

queryVector

The query vector. Ensure that the query vector has the same dimension as the dimension of the vector index that's being queried. For example, if your vector index contains vectors with 384 dimensions, your query vector must also have 384 dimensions.

Type: [VectorData](#) object

Note: This object is a Union. Only one member of this object can be specified or returned.

Required: Yes

returnDistance

Indicates whether to include the computed distance in the response. The default value is `false`.

Type: Boolean

Required: No

returnMetadata

Indicates whether to include metadata in the response. The default value is `false`.

Type: Boolean

Required: No

topK

The number of results to return for each query.

Type: Integer

Valid Range: Minimum value of 1.

Required: Yes

vectorBucketName

The name of the vector bucket that contains the vector index.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "distanceMetric": "string",
  "nextToken": "string",
  "vectors": [
    {
      "distance": number,
      "key": "string",
      "metadata": JSON value
    }
  ]
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[distanceMetric](#)

The distance metric that was used for the similarity search calculation. This is the same distance metric that was configured for the vector index when it was created.

Type: String

Valid Values: euclidean | cosine

[nextToken](#)

Pagination token to be used in the subsequent page request. The field is empty if no further pagination is required.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 4096.

vectors

The vectors in the approximate nearest neighbor search.

Type: Array of [QueryOutputVector](#) objects

Errors

AccessDeniedException

Access denied.

HTTP Status Code: 403

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

KmsDisabledException

The specified AWS KMS key isn't enabled.

HTTP Status Code: 400

KmsInvalidKeyUsageException

The request was rejected for one of the following reasons:

- The KeyUsage value of the KMS key is incompatible with the API operation.
- The encryption algorithm or signing algorithm specified for the operation is incompatible with the type of key material in the KMS key (KeySpec).

For more information, see [InvalidKeyUsageException](#) in the *AWS Key Management Service API Reference*.

HTTP Status Code: 400

KmsInvalidStateException

The key state of the KMS key isn't compatible with the operation.

For more information, see [KMSInvalidStateException](#) in the *AWS Key Management Service API Reference*.

HTTP Status Code: 400

KmsNotFoundException

The KMS key can't be found.

HTTP Status Code: 400

NotFoundException

The request was rejected because the specified resource can't be found.

HTTP Status Code: 404

RequestTimeoutException

The request timed out. Retry your request.

HTTP Status Code: 408

ServiceUnavailableException

The service is unavailable. Wait briefly and retry your request. If it continues to fail, increase your waiting time between retries.

HTTP Status Code: 503

TooManyRequestsException

The request was denied due to request throttling.

HTTP Status Code: 429

ValidationException

The requested action isn't valid.

fieldList

A list of specific validation failures that are encountered during input processing. Each entry in the list contains a path to the field that failed validation and a detailed message that explains why the validation failed. To satisfy multiple constraints, a field can appear multiple times in this list if it failed. You can use the information to identify and fix the specific validation issues in your request.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

TagResource

Service: Amazon S3 Vectors

Applies one or more user-defined tags to an Amazon S3 Vectors resource or updates existing tags. Each tag is a label consisting of a key and value pair. Tags can help you organize, track costs for, and control access to your resources. You can add up to 50 tags for each resource.

Note

For a list of S3 resources that support tagging, see [Managing tags for Amazon S3 resources](#).

Permissions

For vector buckets and vector indexes, you must have the `s3vectors:TagResource` permission to use this operation.

Request Syntax

```
POST /tags/resourceArn HTTP/1.1
Content-type: application/json
```

```
{
  "tags": {
    "string" : "string"
  }
}
```

URI Request Parameters

The request uses the following URI parameters.

resourceArn

The Amazon Resource Name (ARN) of the Amazon S3 Vectors resource that you're applying tags to. The tagged resource can be a vector bucket or a vector index.

Length Constraints: Minimum length of 0. Maximum length of 1011.

Required: Yes

Request Body

The request accepts the following data in JSON format.

tags

The user-defined tag that you want to add to the specified S3 Vectors resource. For more information, see [Tagging for cost allocation or attribute-based access control \(ABAC\)](#).

Type: String to string map

Key Length Constraints: Minimum length of 1. Maximum length of 128.

Key Pattern: ([\p{L}\p{Z}\p{N}_ : / = + \ - @] *)

Value Length Constraints: Minimum length of 0. Maximum length of 256.

Value Pattern: ([\p{L}\p{Z}\p{N}_ : / = + \ - @] *)

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

AccessDeniedException

Access denied.

HTTP Status Code: 403

ConflictException

The request failed because a vector bucket name or a vector index name already exists. Vector bucket names must be unique within your AWS account for each AWS Region. Vector index

names must be unique within your vector bucket. Choose a different vector bucket name or vector index name, and try again.

HTTP Status Code: 409

InternalServerErrorException

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource can't be found.

HTTP Status Code: 404

RequestTimeoutException

The request timed out. Retry your request.

HTTP Status Code: 408

ServiceUnavailableException

The service is unavailable. Wait briefly and retry your request. If it continues to fail, increase your waiting time between retries.

HTTP Status Code: 503

TooManyRequestsException

The request was denied due to request throttling.

HTTP Status Code: 429

ValidationException

The requested action isn't valid.

fieldList

A list of specific validation failures that are encountered during input processing. Each entry in the list contains a path to the field that failed validation and a detailed message that explains why the validation failed. To satisfy multiple constraints, a field can appear multiple times in this list if it failed. You can use the information to identify and fix the specific validation issues in your request.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

UntagResource

Service: Amazon S3 Vectors

Removes the specified user-defined tags from an Amazon S3 Vectors resource. You can pass one or more tag keys.

Note

For a list of S3 resources that support tagging, see [Managing tags for Amazon S3 resources](#).

Permissions

For vector buckets and vector indexes, you must have the `s3vectors:UntagResource` permission to use this operation.

Request Syntax

```
DELETE /tags/resourceArn?tagKeys=tagKeys HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

resourceArn

The Amazon Resource Name (ARN) of the Amazon S3 Vectors resource that you're removing tags from. The tagged resource can be a vector bucket or a vector index.

Length Constraints: Minimum length of 0. Maximum length of 1011.

Required: Yes

tagKeys

The array of tag keys that you're removing from the S3 Vectors resource. For more information, see [Tagging for cost allocation or attribute-based access control \(ABAC\)](#).

Length Constraints: Minimum length of 1. Maximum length of 128.

Pattern: (`[\p{L}\p{Z}\p{N}_.: /+=\ -@]*`)

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

AccessDeniedException

Access denied.

HTTP Status Code: 403

ConflictException

The request failed because a vector bucket name or a vector index name already exists. Vector bucket names must be unique within your AWS account for each AWS Region. Vector index names must be unique within your vector bucket. Choose a different vector bucket name or vector index name, and try again.

HTTP Status Code: 409

InternalServerError

The request failed due to an internal server error.

HTTP Status Code: 500

NotFoundException

The request was rejected because the specified resource can't be found.

HTTP Status Code: 404

RequestTimeoutException

The request timed out. Retry your request.

HTTP Status Code: 408

ServiceUnavailableException

The service is unavailable. Wait briefly and retry your request. If it continues to fail, increase your waiting time between retries.

HTTP Status Code: 503

TooManyRequestsException

The request was denied due to request throttling.

HTTP Status Code: 429

ValidationException

The requested action isn't valid.

fieldList

A list of specific validation failures that are encountered during input processing. Each entry in the list contains a path to the field that failed validation and a detailed message that explains why the validation failed. To satisfy multiple constraints, a field can appear multiple times in this list if it failed. You can use the information to identify and fix the specific validation issues in your request.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

Data Types

The following data types are supported by Amazon S3:

- [AbacStatus](#)
- [AbortIncompleteMultipartUpload](#)
- [AccelerateConfiguration](#)
- [AccessControlPolicy](#)
- [AccessControlTranslation](#)
- [AnalyticsAndOperator](#)
- [AnalyticsConfiguration](#)
- [AnalyticsExportDestination](#)
- [AnalyticsFilter](#)
- [AnalyticsS3BucketDestination](#)
- [AnnotationEntry](#)
- [AnnotationTableConfiguration](#)
- [AnnotationTableConfigurationResult](#)
- [AnnotationTableConfigurationUpdates](#)
- [BlockedEncryptionTypes](#)
- [Bucket](#)
- [BucketInfo](#)
- [BucketLifecycleConfiguration](#)
- [BucketLoggingStatus](#)
- [Checksum](#)
- [CloudFunctionConfiguration](#)
- [CommonPrefix](#)
- [CompletedMultipartUpload](#)
- [CompletedPart](#)
- [Condition](#)
- [ContinuationEvent](#)

- [CopyObjectResult](#)
- [CopyPartResult](#)
- [CORSConfiguration](#)
- [CORSRule](#)
- [CreateBucketConfiguration](#)
- [CSVInput](#)
- [CSVOutput](#)
- [DefaultRetention](#)
- [Delete](#)
- [DeletedObject](#)
- [DeleteMarkerEntry](#)
- [DeleteMarkerReplication](#)
- [Destination](#)
- [DestinationResult](#)
- [Encryption](#)
- [EncryptionConfiguration](#)
- [EndEvent](#)
- [Error](#)
- [ErrorDetails](#)
- [ErrorDocument](#)
- [EventBridgeConfiguration](#)
- [EventHoldDuration](#)
- [ExistingObjectReplication](#)
- [FilterRule](#)
- [GetBucketMetadataConfigurationResult](#)
- [GetBucketMetadataTableConfigurationResult](#)
- [GetObjectAttributesParts](#)
- [GlacierJobParameters](#)
- [Grant](#)
- [Grantee](#)

- [IndexDocument](#)
- [Initiator](#)
- [InputSerialization](#)
- [IntelligentTieringAndOperator](#)
- [IntelligentTieringConfiguration](#)
- [IntelligentTieringFilter](#)
- [InventoryConfiguration](#)
- [InventoryDestination](#)
- [InventoryEncryption](#)
- [InventoryFilter](#)
- [InventoryS3BucketDestination](#)
- [InventorySchedule](#)
- [InventoryTableConfiguration](#)
- [InventoryTableConfigurationResult](#)
- [InventoryTableConfigurationUpdates](#)
- [JournalTableConfiguration](#)
- [JournalTableConfigurationResult](#)
- [JournalTableConfigurationUpdates](#)
- [JSONInput](#)
- [JSONOutput](#)
- [LambdaFunctionConfiguration](#)
- [LifecycleConfiguration](#)
- [LifecycleExpiration](#)
- [LifecycleRule](#)
- [LifecycleRuleAndOperator](#)
- [LifecycleRuleFilter](#)
- [LocationInfo](#)
- [LoggingEnabled](#)
- [MetadataConfiguration](#)
- [MetadataConfigurationResult](#)

- [MetadataEntry](#)
- [MetadataTableConfiguration](#)
- [MetadataTableConfigurationResult](#)
- [MetadataTableEncryptionConfiguration](#)
- [Metrics](#)
- [MetricsAndOperator](#)
- [MetricsConfiguration](#)
- [MetricsFilter](#)
- [MultipartUpload](#)
- [NoncurrentVersionExpiration](#)
- [NoncurrentVersionTransition](#)
- [NotificationConfiguration](#)
- [NotificationConfigurationDeprecated](#)
- [NotificationConfigurationFilter](#)
- [Object](#)
- [ObjectEncryption](#)
- [ObjectIdentifier](#)
- [ObjectLockConfiguration](#)
- [ObjectLockLegalHold](#)
- [ObjectLockRetention](#)
- [ObjectLockRule](#)
- [ObjectPart](#)
- [ObjectVersion](#)
- [OutputLocation](#)
- [OutputSerialization](#)
- [Owner](#)
- [OwnershipControls](#)
- [OwnershipControlsRule](#)
- [ParquetInput](#)
- [Part](#)

- [PartitionedPrefix](#)
- [PolicyStatus](#)
- [Progress](#)
- [ProgressEvent](#)
- [PublicAccessBlockConfiguration](#)
- [QueueConfiguration](#)
- [QueueConfigurationDeprecated](#)
- [RecordExpiration](#)
- [RecordsEvent](#)
- [Redirect](#)
- [RedirectAllRequestsTo](#)
- [ReplicaModifications](#)
- [ReplicationConfiguration](#)
- [ReplicationRule](#)
- [ReplicationRuleAndOperator](#)
- [ReplicationRuleFilter](#)
- [ReplicationTime](#)
- [ReplicationTimeValue](#)
- [RequestPaymentConfiguration](#)
- [RequestProgress](#)
- [RestoreRequest](#)
- [RestoreStatus](#)
- [RoutingRule](#)
- [Rule](#)
- [S3KeyFilter](#)
- [S3Location](#)
- [S3TablesDestination](#)
- [S3TablesDestinationResult](#)
- [ScanRange](#)
- [SelectObjectContentEventStream](#)

- [SelectParameters](#)
- [ServerSideEncryptionByDefault](#)
- [ServerSideEncryptionConfiguration](#)
- [ServerSideEncryptionRule](#)
- [SessionCredentials](#)
- [SimplePrefix](#)
- [SourceSelectionCriteria](#)
- [SSEKMS](#)
- [SseKmsEncryptedObjects](#)
- [SSEKMSEncryption](#)
- [SSES3](#)
- [Stats](#)
- [StatsEvent](#)
- [StorageClassAnalysis](#)
- [StorageClassAnalysisDataExport](#)
- [Tag](#)
- [Tagging](#)
- [TargetGrant](#)
- [TargetObjectKeyFormat](#)
- [Tiering](#)
- [TopicConfiguration](#)
- [TopicConfigurationDeprecated](#)
- [Transition](#)
- [VersioningConfiguration](#)
- [WebsiteConfiguration](#)

The following data types are supported by Amazon S3 Control:

- [AbortIncompleteMultipartUpload](#)
- [AccessControlTranslation](#)
- [AccessGrantsLocationConfiguration](#)

- [AccessPoint](#)
- [AccountLevel](#)
- [ActivityMetrics](#)
- [AdvancedCostOptimizationMetrics](#)
- [AdvancedDataProtectionMetrics](#)
- [AdvancedPerformanceMetrics](#)
- [AsyncErrorDetails](#)
- [AsyncOperation](#)
- [AsyncRequestParameters](#)
- [AsyncResponseDetails](#)
- [AwsLambdaTransformation](#)
- [BucketLevel](#)
- [CloudWatchMetrics](#)
- [CreateBucketConfiguration](#)
- [CreateMultiRegionAccessPointInput](#)
- [Credentials](#)
- [DeleteMarkerReplication](#)
- [DeleteMultiRegionAccessPointInput](#)
- [Destination](#)
- [DetailedStatusCodesMetrics](#)
- [DSSEKMSFilter](#)
- [EncryptionConfiguration](#)
- [EstablishedMultiRegionAccessPointPolicy](#)
- [Exclude](#)
- [ExistingObjectReplication](#)
- [GeneratedManifestEncryption](#)
- [Grantee](#)
- [Include](#)
- [JobDescriptor](#)
- [JobFailure](#)

- [JobListDescriptor](#)
- [JobManifest](#)
- [JobManifestGenerator](#)
- [JobManifestGeneratorFilter](#)
- [JobManifestLocation](#)
- [JobManifestSpec](#)
- [JobOperation](#)
- [JobProgressSummary](#)
- [JobReport](#)
- [JobTimers](#)
- [KeyNameConstraint](#)
- [LambdaInvokeOperation](#)
- [LifecycleConfiguration](#)
- [LifecycleExpiration](#)
- [LifecycleRule](#)
- [LifecycleRuleAndOperator](#)
- [LifecycleRuleFilter](#)
- [ListAccessGrantEntry](#)
- [ListAccessGrantsInstanceEntry](#)
- [ListAccessGrantsLocationsEntry](#)
- [ListCallerAccessGrantsEntry](#)
- [ListStorageLensConfigurationEntry](#)
- [ListStorageLensGroupEntry](#)
- [MatchObjectAge](#)
- [MatchObjectSize](#)
- [Metrics](#)
- [MultiRegionAccessPointPolicyDocument](#)
- [MultiRegionAccessPointRegionalResponse](#)
- [MultiRegionAccessPointReport](#)
- [MultiRegionAccessPointRoute](#)

- [MultiRegionAccessPointsAsyncResponse](#)
- [NoncurrentVersionExpiration](#)
- [NoncurrentVersionTransition](#)
- [NotSSEFilter](#)
- [ObjectEncryption](#)
- [ObjectEncryptionFilter](#)
- [ObjectLambdaAccessPoint](#)
- [ObjectLambdaAccessPointAlias](#)
- [ObjectLambdaConfiguration](#)
- [ObjectLambdaContentTransformation](#)
- [ObjectLambdaTransformationConfiguration](#)
- [PolicyStatus](#)
- [PrefixLevel](#)
- [PrefixLevelStorageMetrics](#)
- [ProposedMultiRegionAccessPointPolicy](#)
- [PublicAccessBlockConfiguration](#)
- [PutMultiRegionAccessPointPolicyInput](#)
- [Region](#)
- [RegionalBucket](#)
- [RegionReport](#)
- [ReplicaModifications](#)
- [ReplicationConfiguration](#)
- [ReplicationRule](#)
- [ReplicationRuleAndOperator](#)
- [ReplicationRuleFilter](#)
- [ReplicationTime](#)
- [ReplicationTimeValue](#)
- [S3AccessControlList](#)
- [S3AccessControlPolicy](#)
- [S3BucketDestination](#)

- [S3ComputeObjectChecksumOperation](#)
- [S3CopyObjectOperation](#)
- [S3DeleteObjectTaggingOperation](#)
- [S3GeneratedManifestDescriptor](#)
- [S3Grant](#)
- [S3Grantee](#)
- [S3InitiateRestoreObjectOperation](#)
- [S3JobManifestGenerator](#)
- [S3ManifestOutputLocation](#)
- [S3ObjectLockEventHoldDuration](#)
- [S3ObjectLockLegalHold](#)
- [S3ObjectLockRetentionEventHoldDuration](#)
- [S3ObjectMetadata](#)
- [S3ObjectOwner](#)
- [S3ReplicateObjectOperation](#)
- [S3Retention](#)
- [S3SetObjectAclOperation](#)
- [S3SetObjectLegalHoldOperation](#)
- [S3SetObjectRetentionOperation](#)
- [S3SetObjectTaggingOperation](#)
- [S3Tag](#)
- [S3UpdateObjectEncryptionOperation](#)
- [S3UpdateObjectEncryptionSSEKMS](#)
- [Scope](#)
- [SelectionCriteria](#)
- [SourceSelectionCriteria](#)
- [SSECFilter](#)
- [SSEKMS](#)
- [SseKmsEncryptedObjects](#)
- [SSEKMSEncryption](#)

- [SSEKMSFilter](#)
- [SSES3](#)
- [SSES3Encryption](#)
- [SSES3Filter](#)
- [StorageLensAwsOrg](#)
- [StorageLensConfiguration](#)
- [StorageLensDataExport](#)
- [StorageLensDataExportEncryption](#)
- [StorageLensExpandedPrefixesDataExport](#)
- [StorageLensGroup](#)
- [StorageLensGroupAndOperator](#)
- [StorageLensGroupFilter](#)
- [StorageLensGroupLevel](#)
- [StorageLensGroupLevelSelectionCriteria](#)
- [StorageLensGroupOrOperator](#)
- [StorageLensTableDestination](#)
- [StorageLensTag](#)
- [Tag](#)
- [Tagging](#)
- [Transition](#)
- [VersioningConfiguration](#)
- [VpcConfiguration](#)

The following data types are supported by Amazon S3 Files:

- [CreationPermissions](#)
- [ExpirationDataRule](#)
- [ImportDataRule](#)
- [ListAccessPointsDescription](#)
- [ListFileSystemsDescription](#)
- [ListMountTargetsDescription](#)

- [PosixUser](#)
- [RootDirectory](#)
- [Tag](#)

The following data types are supported by Amazon S3 on Outposts:

- [Endpoint](#)
- [FailedReason](#)
- [NetworkInterface](#)
- [Outpost](#)

The following data types are supported by Amazon S3 Tables:

- [EncryptionConfiguration](#)
- [IcebergCompactionSettings](#)
- [IcebergMetadata](#)
- [IcebergPartitionField](#)
- [IcebergPartitionSpec](#)
- [IcebergSchema](#)
- [IcebergSchemaV2](#)
- [IcebergSnapshotManagementSettings](#)
- [IcebergSortField](#)
- [IcebergSortOrder](#)
- [IcebergUnreferencedFileRemovalSettings](#)
- [LastSuccessfulReplicatedUpdate](#)
- [ManagedTableInformation](#)
- [NamespaceSummary](#)
- [ReplicationDestination](#)
- [ReplicationDestinationStatusModel](#)
- [ReplicationInformation](#)
- [SchemaField](#)
- [SchemaV2Field](#)

- [StorageClassConfiguration](#)
- [TableBucketMaintenanceConfigurationValue](#)
- [TableBucketMaintenanceSettings](#)
- [TableBucketReplicationConfiguration](#)
- [TableBucketReplicationRule](#)
- [TableBucketSummary](#)
- [TableMaintenanceConfigurationValue](#)
- [TableMaintenanceJobStatusValue](#)
- [TableMaintenanceSettings](#)
- [TableMetadata](#)
- [TableRecordExpirationConfigurationValue](#)
- [TableRecordExpirationJobMetrics](#)
- [TableRecordExpirationSettings](#)
- [TableReplicationConfiguration](#)
- [TableReplicationRule](#)
- [TableSummary](#)

The following data types are supported by Amazon S3 Vectors:

- [EncryptionConfiguration](#)
- [GetOutputVector](#)
- [Index](#)
- [IndexSummary](#)
- [ListOutputVector](#)
- [MetadataConfiguration](#)
- [PutInputVector](#)
- [QueryOutputVector](#)
- [ValidationExceptionField](#)
- [VectorBucket](#)
- [VectorBucketSummary](#)
- [VectorData](#)

Amazon S3

The following data types are supported by Amazon S3:

- [AbacStatus](#)
- [AbortIncompleteMultipartUpload](#)
- [AccelerateConfiguration](#)
- [AccessControlPolicy](#)
- [AccessControlTranslation](#)
- [AnalyticsAndOperator](#)
- [AnalyticsConfiguration](#)
- [AnalyticsExportDestination](#)
- [AnalyticsFilter](#)
- [AnalyticsS3BucketDestination](#)
- [AnnotationEntry](#)
- [AnnotationTableConfiguration](#)
- [AnnotationTableConfigurationResult](#)
- [AnnotationTableConfigurationUpdates](#)
- [BlockedEncryptionTypes](#)
- [Bucket](#)
- [BucketInfo](#)
- [BucketLifecycleConfiguration](#)
- [BucketLoggingStatus](#)
- [Checksum](#)
- [CloudFunctionConfiguration](#)
- [CommonPrefix](#)
- [CompletedMultipartUpload](#)
- [CompletedPart](#)
- [Condition](#)
- [ContinuationEvent](#)
- [CopyObjectResult](#)

- [CopyPartResult](#)
- [CORSConfiguration](#)
- [CORSRule](#)
- [CreateBucketConfiguration](#)
- [CSVInput](#)
- [CSVOutput](#)
- [DefaultRetention](#)
- [Delete](#)
- [DeletedObject](#)
- [DeleteMarkerEntry](#)
- [DeleteMarkerReplication](#)
- [Destination](#)
- [DestinationResult](#)
- [Encryption](#)
- [EncryptionConfiguration](#)
- [EndEvent](#)
- [Error](#)
- [ErrorDetails](#)
- [ErrorDocument](#)
- [EventBridgeConfiguration](#)
- [EventHoldDuration](#)
- [ExistingObjectReplication](#)
- [FilterRule](#)
- [GetBucketMetadataConfigurationResult](#)
- [GetBucketMetadataTableConfigurationResult](#)
- [GetObjectAttributesParts](#)
- [GlacierJobParameters](#)
- [Grant](#)
- [Grantee](#)
- [IndexDocument](#)

- [Initiator](#)
- [InputSerialization](#)
- [IntelligentTieringAndOperator](#)
- [IntelligentTieringConfiguration](#)
- [IntelligentTieringFilter](#)
- [InventoryConfiguration](#)
- [InventoryDestination](#)
- [InventoryEncryption](#)
- [InventoryFilter](#)
- [InventoryS3BucketDestination](#)
- [InventorySchedule](#)
- [InventoryTableConfiguration](#)
- [InventoryTableConfigurationResult](#)
- [InventoryTableConfigurationUpdates](#)
- [JournalTableConfiguration](#)
- [JournalTableConfigurationResult](#)
- [JournalTableConfigurationUpdates](#)
- [JSONInput](#)
- [JSONOutput](#)
- [LambdaFunctionConfiguration](#)
- [LifecycleConfiguration](#)
- [LifecycleExpiration](#)
- [LifecycleRule](#)
- [LifecycleRuleAndOperator](#)
- [LifecycleRuleFilter](#)
- [LocationInfo](#)
- [LoggingEnabled](#)
- [MetadataConfiguration](#)
- [MetadataConfigurationResult](#)
- [MetadataEntry](#)

- [MetadataTableConfiguration](#)
- [MetadataTableConfigurationResult](#)
- [MetadataTableEncryptionConfiguration](#)
- [Metrics](#)
- [MetricsAndOperator](#)
- [MetricsConfiguration](#)
- [MetricsFilter](#)
- [MultipartUpload](#)
- [NoncurrentVersionExpiration](#)
- [NoncurrentVersionTransition](#)
- [NotificationConfiguration](#)
- [NotificationConfigurationDeprecated](#)
- [NotificationConfigurationFilter](#)
- [Object](#)
- [ObjectEncryption](#)
- [ObjectIdentifier](#)
- [ObjectLockConfiguration](#)
- [ObjectLockLegalHold](#)
- [ObjectLockRetention](#)
- [ObjectLockRule](#)
- [ObjectPart](#)
- [ObjectVersion](#)
- [OutputLocation](#)
- [OutputSerialization](#)
- [Owner](#)
- [OwnershipControls](#)
- [OwnershipControlsRule](#)
- [ParquetInput](#)
- [Part](#)
- [PartitionedPrefix](#)

- [PolicyStatus](#)
- [Progress](#)
- [ProgressEvent](#)
- [PublicAccessBlockConfiguration](#)
- [QueueConfiguration](#)
- [QueueConfigurationDeprecated](#)
- [RecordExpiration](#)
- [RecordsEvent](#)
- [Redirect](#)
- [RedirectAllRequestsTo](#)
- [ReplicaModifications](#)
- [ReplicationConfiguration](#)
- [ReplicationRule](#)
- [ReplicationRuleAndOperator](#)
- [ReplicationRuleFilter](#)
- [ReplicationTime](#)
- [ReplicationTimeValue](#)
- [RequestPaymentConfiguration](#)
- [RequestProgress](#)
- [RestoreRequest](#)
- [RestoreStatus](#)
- [RoutingRule](#)
- [Rule](#)
- [S3KeyFilter](#)
- [S3Location](#)
- [S3TablesDestination](#)
- [S3TablesDestinationResult](#)
- [ScanRange](#)
- [SelectObjectContentEventStream](#)
- [SelectParameters](#)

- [ServerSideEncryptionByDefault](#)
- [ServerSideEncryptionConfiguration](#)
- [ServerSideEncryptionRule](#)
- [SessionCredentials](#)
- [SimplePrefix](#)
- [SourceSelectionCriteria](#)
- [SSEKMS](#)
- [SseKmsEncryptedObjects](#)
- [SSEKMSEncryption](#)
- [SSES3](#)
- [Stats](#)
- [StatsEvent](#)
- [StorageClassAnalysis](#)
- [StorageClassAnalysisDataExport](#)
- [Tag](#)
- [Tagging](#)
- [TargetGrant](#)
- [TargetObjectKeyFormat](#)
- [Tiering](#)
- [TopicConfiguration](#)
- [TopicConfigurationDeprecated](#)
- [Transition](#)
- [VersioningConfiguration](#)
- [WebsiteConfiguration](#)

AbacStatus

Service: Amazon S3

The ABAC status of the general purpose bucket. When ABAC is enabled for the general purpose bucket, you can use tags to manage access to the general purpose buckets as well as for cost tracking purposes. When ABAC is disabled for the general purpose buckets, you can only use tags for cost tracking purposes. For more information, see [Using tags with S3 general purpose buckets](#).

Contents

Status

The ABAC status of the general purpose bucket.

Type: String

Valid Values: Enabled | Disabled

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

AbortIncompleteMultipartUpload

Service: Amazon S3

Specifies the days since the initiation of an incomplete multipart upload that Amazon S3 will wait before permanently removing all parts of the upload. For more information, see [Aborting Incomplete Multipart Uploads Using a Bucket Lifecycle Configuration](#) in the *Amazon S3 User Guide*.

Contents

DaysAfterInitiation

Specifies the number of days after which Amazon S3 aborts an incomplete multipart upload.

Type: Integer

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

AccelerateConfiguration

Service: Amazon S3

Configures the transfer acceleration state for an Amazon S3 bucket. For more information, see [Amazon S3 Transfer Acceleration](#) in the *Amazon S3 User Guide*.

Contents

Status

Specifies the transfer acceleration status of the bucket.

Type: String

Valid Values: Enabled | Suspended

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

AccessControlPolicy

Service: Amazon S3

Contains the elements that set the ACL permissions for an object per grantee.

Contents

Grants

A list of grants.

Type: Array of [Grant](#) data types

Required: No

Owner

Container for the bucket owner's display name and ID.

Type: [Owner](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

AccessControlTranslation

Service: Amazon S3

A container for information about access control for replicas.

Contents

Owner

Specifies the replica ownership. For default and valid values, see [PUT bucket replication](#) in the *Amazon S3 API Reference*.

Type: String

Valid Values: Destination

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

AnalyticsAndOperator

Service: Amazon S3

A conjunction (logical AND) of predicates, which is used in evaluating a metrics filter. The operator must have at least two predicates in any combination, and an object must match all of the predicates for the filter to apply.

Contents

Prefix

The prefix to use when evaluating an AND predicate: The prefix that an object must have to be included in the metrics results.

Type: String

Required: No

Tags

The list of tags to use when evaluating an AND predicate.

Type: Array of [Tag](#) data types

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

AnalyticsConfiguration

Service: Amazon S3

Specifies the configuration and any analyses for the analytics filter of an Amazon S3 bucket.

Contents

Id

The ID that identifies the analytics configuration.

Type: String

Required: Yes

StorageClassAnalysis

Contains data related to access patterns to be collected and made available to analyze the tradeoffs between different storage classes.

Type: [StorageClassAnalysis](#) data type

Required: Yes

Filter

The filter used to describe a set of objects for analyses. A filter must have exactly one prefix, one tag, or one conjunction (AnalyticsAndOperator). If no filter is provided, all objects will be considered in any analysis.

Type: [AnalyticsFilter](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

AnalyticsExportDestination

Service: Amazon S3

Where to publish the analytics results.

Contents

S3BucketDestination

A destination signifying output to an S3 bucket.

Type: [AnalyticsS3BucketDestination](#) data type

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

AnalyticsFilter

Service: Amazon S3

The filter used to describe a set of objects for analyses. A filter must have exactly one prefix, one tag, or one conjunction (AnalyticsAndOperator). If no filter is provided, all objects will be considered in any analysis.

Contents

And

A conjunction (logical AND) of predicates, which is used in evaluating an analytics filter. The operator must have at least two predicates.

Type: [AnalyticsAndOperator](#) data type

Required: No

Prefix

The prefix to use when evaluating an analytics filter.

Type: String

Required: No

Tag

The tag to use when evaluating an analytics filter.

Type: [Tag](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

AnalyticsS3BucketDestination

Service: Amazon S3

Contains information about where to publish the analytics results.

Contents

Bucket

The Amazon Resource Name (ARN) of the bucket to which data is exported.

Type: String

Required: Yes

Format

Specifies the file format used when exporting data to Amazon S3.

Type: String

Valid Values: CSV

Required: Yes

BucketAccountId

The account ID that owns the destination S3 bucket. If no account ID is provided, the owner is not validated before exporting data.

Note

Although this value is optional, we strongly recommend that you set it to help prevent problems if the destination bucket ownership changes.

Type: String

Required: No

Prefix

The prefix to use when exporting data. The prefix is prepended to all results.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

AnnotationEntry

Service: Amazon S3

Describes a single annotation attached to an object, including its name, last modified time, size, ETag, checksum algorithm, and replication status. Returned in the response from `ListObjectAnnotations`.

Contents

AnnotationName

The name of the annotation.

Type: String

Required: Yes

LastModified

The date and time the annotation was last modified.

Type: Timestamp

Required: Yes

Size

The size of the annotation payload, in bytes.

Type: Long

Required: Yes

ChecksumAlgorithm

The checksum algorithm used for the annotation.

Type: Array of strings

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 |
XXHASH64 | XXHASH3 | XXHASH128

Required: No

ETag

The entity tag of the annotation.

Type: String

Required: No

ReplicationStatus

The replication status of the annotation.

Type: String

Valid Values: COMPLETE | PENDING | FAILED | REPLICA | COMPLETED

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

AnnotationTableConfiguration

Service: Amazon S3

Specifies the configuration for the annotation table associated with a bucket's Amazon S3 Metadata configuration. The annotation table is an Iceberg table that records annotation events for objects in the bucket.

Contents

ConfigurationState

The state of the annotation table. Valid values are ENABLED and DISABLED.

Type: String

Valid Values: ENABLED | DISABLED

Required: Yes

EncryptionConfiguration

The encryption settings for an S3 Metadata journal table or inventory table configuration.

Type: [MetadataTableEncryptionConfiguration](#) data type

Required: No

Role

The ARN of the IAM role used to manage the annotation table.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)

- [AWS SDK for Ruby V3](#)

AnnotationTableConfigurationResult

Service: Amazon S3

Contains the current state of the annotation table associated with a bucket's Amazon S3 Metadata configuration, including its provisioning status and identifiers.

Contents

ConfigurationState

The current configuration state of the annotation table.

Type: String

Valid Values: ENABLED | DISABLED

Required: Yes

Error

If an S3 Metadata V1 `CreateBucketMetadataTableConfiguration` or V2 `CreateBucketMetadataConfiguration` request succeeds, but S3 Metadata was unable to create the table, this structure contains the error code and error message.

Note

If you created your S3 Metadata configuration before July 15, 2025, we recommend that you delete and re-create your configuration by using [CreateBucketMetadataConfiguration](#) so that you can expire journal table records and create a live inventory table.

Type: [ErrorDetails](#) data type

Required: No

Role

The ARN of the IAM role associated with the annotation table.

Type: String

Required: No

TableArn

The ARN of the annotation table.

Type: String

Required: No

TableName

The name of the annotation table.

Type: String

Required: No

TableStatus

The provisioning status of the annotation table. Possible values: CREATING, BACKFILLING, ACTIVE, FAILED.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

AnnotationTableConfigurationUpdates

Service: Amazon S3

Specifies updates to apply to the annotation table configuration. Used as the request body for `UpdateBucketMetadataAnnotationTableConfiguration`.

Contents

ConfigurationState

The new configuration state to apply.

Type: String

Valid Values: ENABLED | DISABLED

Required: Yes

EncryptionConfiguration

The encryption settings for an S3 Metadata journal table or inventory table configuration.

Type: [MetadataTableEncryptionConfiguration](#) data type

Required: No

Role

The new IAM role ARN to apply.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

BlockedEncryptionTypes

Service: Amazon S3

A bucket-level setting for Amazon S3 general purpose buckets used to prevent the upload of new objects encrypted with the specified server-side encryption type. For example, blocking an encryption type will block PutObject, CopyObject, PostObject, multipart upload, and replication requests to the bucket for objects with the specified encryption type. However, you can continue to read and list any pre-existing objects already encrypted with the specified encryption type. For more information, see [Blocking or unblocking SSE-C for a general purpose bucket](#).

This data type is used with the following actions:

- [PutBucketEncryption](#)
- [GetBucketEncryption](#)
- [DeleteBucketEncryption](#)

Permissions

You must have the `s3:PutEncryptionConfiguration` permission to block or unblock an encryption type for a bucket.

You must have the `s3:GetEncryptionConfiguration` permission to view a bucket's encryption type.

Contents

EncryptionType

The object encryption type that you want to block or unblock for an Amazon S3 general purpose bucket.

Note

Currently, this parameter only supports blocking or unblocking server side encryption with customer-provided keys (SSE-C). For more information about SSE-C, see [Using server-side encryption with customer-provided keys \(SSE-C\)](#).

Type: Array of strings

Valid Values: NONE | SSE-C

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Bucket

Service: Amazon S3

In terms of implementation, a Bucket is a resource.

Contents

BucketArn

The Amazon Resource Name (ARN) of the S3 bucket. ARNs uniquely identify AWS resources across all of AWS.

Note

This parameter is only supported for S3 directory buckets. For more information, see [Using tags with directory buckets](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 128.

Pattern: `arn:[^:]+:(s3|s3express):.*`

Required: No

BucketRegion

BucketRegion indicates the AWS region where the bucket is located. If the request contains at least one valid parameter, it is included in the response.

Type: String

Required: No

CreationDate

Date the bucket was created. This date can change when making changes to your bucket, such as editing its bucket policy.

Type: Timestamp

Required: No

Name

The name of the bucket.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

BucketInfo

Service: Amazon S3

Specifies the information about the bucket that will be created. For more information about directory buckets, see [Directory buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is only supported by directory buckets.

Contents

DataRedundancy

The number of Zone (Availability Zone or Local Zone) that's used for redundancy for the bucket.

Type: String

Valid Values: `SingleAvailabilityZone` | `SingleLocalZone`

Required: No

Type

The type of bucket.

Type: String

Valid Values: `Directory`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

BucketLifecycleConfiguration

Service: Amazon S3

Specifies the lifecycle configuration for objects in an Amazon S3 bucket. For more information, see [Object Lifecycle Management](#) in the *Amazon S3 User Guide*.

Contents

Rules

A lifecycle rule for individual objects in an Amazon S3 bucket.

Type: Array of [LifecycleRule](#) data types

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

BucketLoggingStatus

Service: Amazon S3

Container for logging status information.

Contents

LoggingEnabled

Describes where logs are stored and the prefix that Amazon S3 assigns to all log object keys for a bucket. For more information, see [PUT Bucket logging](#) in the *Amazon S3 API Reference*.

Type: [LoggingEnabled](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Checksum

Service: Amazon S3

Contains all the possible checksum or digest values for an object.

Contents

ChecksumCRC32

The Base64 encoded, 32-bit CRC32 checksum of the object. This checksum is only present if the checksum was uploaded with the object. When you use an API operation on an object that was uploaded using multipart uploads, this value may not be a direct checksum value of the full object. Instead, it's a calculation based on the checksum values of each individual part. For more information about how checksums are calculated with multipart uploads, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumCRC32C

The Base64 encoded, 32-bit CRC32C checksum of the object. This checksum is only present if the checksum was uploaded with the object. When you use an API operation on an object that was uploaded using multipart uploads, this value may not be a direct checksum value of the full object. Instead, it's a calculation based on the checksum values of each individual part. For more information about how checksums are calculated with multipart uploads, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumCRC64NVME

The Base64 encoded, 64-bit CRC64NVME checksum of the object. This checksum is present if the object was uploaded with the CRC64NVME checksum algorithm, or if the object was uploaded without a checksum (and Amazon S3 added the default checksum, CRC64NVME, to the uploaded object). For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumMD5

The Base64 encoded, 128-bit MD5 digest of the object. This checksum is present if the object was uploaded with the MD5 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumSHA1

The Base64 encoded, 160-bit SHA1 digest of the object. This checksum is only present if the checksum was uploaded with the object. When you use the API operation on an object that was uploaded using multipart uploads, this value may not be a direct checksum value of the full object. Instead, it's a calculation based on the checksum values of each individual part. For more information about how checksums are calculated with multipart uploads, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumSHA256

The Base64 encoded, 256-bit SHA256 digest of the object. This checksum is only present if the checksum was uploaded with the object. When you use an API operation on an object that was uploaded using multipart uploads, this value may not be a direct checksum value of the full object. Instead, it's a calculation based on the checksum values of each individual part. For more information about how checksums are calculated with multipart uploads, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumSHA512

The Base64 encoded, 512-bit SHA512 digest of the object. This checksum is present if the object was uploaded with the SHA512 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumType

The checksum type that is used to calculate the object's checksum value. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Valid Values: COMPOSITE | FULL_OBJECT

Required: No

ChecksumXXHASH128

The Base64 encoded, 128-bit XXHASH128 checksum of the object. This checksum is present if the object was uploaded with the XXHASH128 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumXXHASH3

The Base64 encoded, 64-bit XXHASH3 checksum of the object. This checksum is present if the object was uploaded with the XXHASH3 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumXXHASH64

The Base64 encoded, 64-bit XXHASH64 checksum of the object. This checksum is present if the object was uploaded with the XXHASH64 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

CloudFunctionConfiguration

Service: Amazon S3

Container for specifying the AWS Lambda notification configuration.

Contents

CloudFunction

Lambda cloud function ARN that Amazon S3 can invoke when it detects events of the specified type.

Type: String

Required: No

Event

This member has been deprecated.

The bucket event for which to send notifications.

Type: String

Valid Values: s3:ReducedRedundancyLostObject | s3:ObjectCreated:* | s3:ObjectCreated:Put | s3:ObjectCreated:Post | s3:ObjectCreated:Copy | s3:ObjectCreated:CompleteMultipartUpload | s3:ObjectRemoved:* | s3:ObjectRemoved:Delete | s3:ObjectRemoved:DeleteMarkerCreated | s3:ObjectRestore:* | s3:ObjectRestore:Post | s3:ObjectRestore:Completed | s3:Replication:* | s3:Replication:OperationFailedReplication | s3:Replication:OperationNotTracked | s3:Replication:OperationMissedThreshold | s3:Replication:OperationReplicatedAfterThreshold | s3:ObjectRestore:Delete | s3:LifecycleTransition | s3: IntelligentTiering | s3:ObjectAcl:Put | s3:LifecycleExpiration:* | s3:LifecycleExpiration:Delete | s3:LifecycleExpiration:DeleteMarkerCreated | s3:ObjectTagging:* | s3:ObjectTagging:Put | s3:ObjectTagging:Delete | s3:ObjectAnnotation:* | s3:ObjectAnnotation:Put | s3:ObjectAnnotation:Delete | s3:ObjectRetention:Put

Required: No

Events

Bucket events for which to send notifications.

Type: Array of strings

Valid Values: s3:ReducedRedundancyLostObject | s3:ObjectCreated:* | s3:ObjectCreated:Put | s3:ObjectCreated:Post | s3:ObjectCreated:Copy | s3:ObjectCreated:CompleteMultipartUpload | s3:ObjectRemoved:* | s3:ObjectRemoved:Delete | s3:ObjectRemoved:DeleteMarkerCreated | s3:ObjectRestore:* | s3:ObjectRestore:Post | s3:ObjectRestore:Completed | s3:Replication:* | s3:Replication:OperationFailedReplication | s3:Replication:OperationNotTracked | s3:Replication:OperationMissedThreshold | s3:Replication:OperationReplicatedAfterThreshold | s3:ObjectRestore:Delete | s3:LifecycleTransition | s3:IntelligentTiering | s3:ObjectAcl:Put | s3:LifecycleExpiration:* | s3:LifecycleExpiration:Delete | s3:LifecycleExpiration:DeleteMarkerCreated | s3:ObjectTagging:* | s3:ObjectTagging:Put | s3:ObjectTagging:Delete | s3:ObjectAnnotation:* | s3:ObjectAnnotation:Put | s3:ObjectAnnotation:Delete | s3:ObjectRetention:Put

Required: No

Id

An optional unique identifier for configurations in a notification configuration. If you don't provide one, Amazon S3 will assign an ID.

Type: String

Required: No

InvocationRole

The role supporting the invocation of the Lambda function

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

CommonPrefix

Service: Amazon S3

Container for all (if there are any) keys between Prefix and the next occurrence of the string specified by a delimiter. CommonPrefixes lists keys that act like subdirectories in the directory specified by Prefix. For example, if the prefix is notes/ and the delimiter is a slash (/) as in notes/summer/july, the common prefix is notes/summer/.

Contents

Prefix

Container for the specified common prefix.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

CompletedMultipartUpload

Service: Amazon S3

The container for the completed multipart upload details.

Contents

Parts

Array of CompletedPart data types.

If you do not supply a valid Part with your request, the service sends back an HTTP 400 response.

Type: Array of [CompletedPart](#) data types

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

CompletedPart

Service: Amazon S3

Details of the parts that were uploaded.

Contents

ChecksumCRC32

The Base64 encoded, 32-bit CRC32 checksum of the part. This checksum is present if the multipart upload request was created with the CRC32 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumCRC32C

The Base64 encoded, 32-bit CRC32C checksum of the part. This checksum is present if the multipart upload request was created with the CRC32C checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumCRC64NVME

The Base64 encoded, 64-bit CRC64NVME checksum of the part. This checksum is present if the multipart upload request was created with the CRC64NVME checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumMD5

The Base64 encoded, 128-bit MD5 digest of the part. This checksum is present if the multipart upload request was created with the MD5 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumSHA1

The Base64 encoded, 160-bit SHA1 checksum of the part. This checksum is present if the multipart upload request was created with the SHA1 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumSHA256

The Base64 encoded, 256-bit SHA256 checksum of the part. This checksum is present if the multipart upload request was created with the SHA256 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumSHA512

The Base64 encoded, 512-bit SHA512 digest of the part. This checksum is present if the multipart upload request was created with the SHA512 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumXXHASH128

The Base64 encoded, 128-bit XXHASH128 checksum of the part. This checksum is present if the multipart upload request was created with the XXHASH128 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumXXHASH3

The Base64 encoded, 64-bit XXHASH3 checksum of the part. This checksum is present if the multipart upload request was created with the XXHASH3 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumXXHASH64

The Base64 encoded, 64-bit XXHASH64 checksum of the part. This checksum is present if the multipart upload request was created with the XXHASH64 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ETag

Entity tag returned when the part was uploaded.

Type: String

Required: No

PartNumber

Part number that identifies the part. This is a positive integer between 1 and 10,000.

Note

- **General purpose buckets** - In `CompleteMultipartUpload`, when a additional checksum (including `x-amz-checksum-crc32`, `x-amz-checksum-crc32c`, `x-amz-checksum-sha1`, or `x-amz-checksum-sha256`) is applied to each part, the `PartNumber` must start at 1 and the part numbers must be consecutive. Otherwise, Amazon S3 generates an HTTP 400 Bad Request status code and an `InvalidPartOrder` error code.
- **Directory buckets** - In `CompleteMultipartUpload`, the `PartNumber` must start at 1 and the part numbers must be consecutive.

Type: Integer

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Condition

Service: Amazon S3

A container for describing a condition that must be met for the specified redirect to apply. For example, 1. If request is for pages in the /docs folder, redirect to the /documents folder. 2. If request results in HTTP error 4xx, redirect request to another host where you might process the error.

Contents

HttpErrorCodeReturnedEquals

The HTTP error code when the redirect is applied. In the event of an error, if the error code equals this value, then the specified redirect is applied. Required when parent element `Condition` is specified and sibling `KeyPrefixEquals` is not specified. If both are specified, then both must be true for the redirect to be applied.

Type: String

Required: No

KeyPrefixEquals

The object key name prefix when the redirect is applied. For example, to redirect requests for `ExamplePage.html`, the key prefix will be `ExamplePage.html`. To redirect request for all pages with the prefix `docs/`, the key prefix will be `/docs`, which identifies all objects in the `docs/` folder. Required when the parent element `Condition` is specified and sibling `HttpErrorCodeReturnedEquals` is not specified. If both conditions are specified, both must be true for the redirect to be applied.

Important

Replacement must be made for object keys containing special characters (such as carriage returns) when using XML requests. For more information, see [XML related object key constraints](#).

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ContinuationEvent

Service: Amazon S3

Contents

The members of this exception structure are context-dependent.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

CopyObjectResult

Service: Amazon S3

Container for all response elements.

Contents

ChecksumCRC32

The Base64 encoded, 32-bit CRC32 checksum of the object. This checksum is only present if the checksum was uploaded with the object. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumCRC32C

The Base64 encoded, 32-bit CRC32C checksum of the object. This checksum is only present if the checksum was uploaded with the object. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumCRC64NVME

The Base64 encoded, 64-bit CRC64NVME checksum of the object. This checksum is present if the object being copied was uploaded with the CRC64NVME checksum algorithm, or if the object was uploaded without a checksum (and Amazon S3 added the default checksum, CRC64NVME, to the uploaded object). For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumMD5

The Base64 encoded, 128-bit MD5 digest of the object. This checksum is only present if the object was uploaded with the MD5 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumSHA1

The Base64 encoded, 160-bit SHA1 digest of the object. This checksum is only present if the checksum was uploaded with the object. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumSHA256

The Base64 encoded, 256-bit SHA256 digest of the object. This checksum is only present if the checksum was uploaded with the object. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumSHA512

The Base64 encoded, 512-bit SHA512 digest of the object. This checksum is only present if the object was uploaded with the SHA512 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumType

The checksum type that is used to calculate the object's checksum value. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Valid Values: COMPOSITE | FULL_OBJECT

Required: No

ChecksumXXHASH128

The Base64 encoded, 128-bit XXHASH128 checksum of the object. This checksum is only present if the object was uploaded with the XXHASH128 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumXXHASH3

The Base64 encoded, 64-bit XXHASH3 checksum of the object. This checksum is only present if the object was uploaded with the XXHASH3 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumXXHASH64

The Base64 encoded, 64-bit XXHASH64 checksum of the object. This checksum is only present if the object was uploaded with the XXHASH64 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ETag

Returns the ETag of the new object. The ETag reflects only changes to the contents of an object, not its metadata.

Type: String

Required: No

LastModified

Creation date of the object.

Type: Timestamp

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

CopyPartResult

Service: Amazon S3

Container for all response elements.

Contents

ChecksumCRC32

The Base64 encoded, 32-bit CRC32 checksum of the part. This checksum is present if the multipart upload request was created with the CRC32 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumCRC32C

The Base64 encoded, 32-bit CRC32C checksum of the part. This checksum is present if the multipart upload request was created with the CRC32C checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumCRC64NVME

The Base64 encoded, 64-bit CRC64NVME checksum of the part. This checksum is present if the multipart upload request was created with the CRC64NVME checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumMD5

The Base64 encoded, 128-bit MD5 digest of the part. This checksum is present if the multipart upload request was created with the MD5 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumSHA1

The Base64 encoded, 160-bit SHA1 digest of the part. This checksum is present if the multipart upload request was created with the SHA1 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumSHA256

The Base64 encoded, 256-bit SHA256 digest of the part. This checksum is present if the multipart upload request was created with the SHA256 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumSHA512

The Base64 encoded, 512-bit SHA512 digest of the part. This checksum is present if the multipart upload request was created with the SHA512 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumXXHASH128

The Base64 encoded, 128-bit XXHASH128 checksum of the part. This checksum is present if the multipart upload request was created with the XXHASH128 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumXXHASH3

The Base64 encoded, 64-bit XXHASH3 checksum of the part. This checksum is present if the multipart upload request was created with the XXHASH3 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumXXHASH64

The Base64 encoded, 64-bit XXHASH64 checksum of the part. This checksum is present if the multipart upload request was created with the XXHASH64 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ETag

Entity tag of the object.

Type: String

Required: No

LastModified

Date and time at which the object was uploaded.

Type: Timestamp

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)

- [AWS SDK for Ruby V3](#)

CORSConfiguration

Service: Amazon S3

Describes the cross-origin access configuration for objects in an Amazon S3 bucket. For more information, see [Enabling Cross-Origin Resource Sharing](#) in the *Amazon S3 User Guide*.

Contents

CORSRules

A set of origins and methods (cross-origin access that you want to allow). You can add up to 100 rules to the configuration.

Type: Array of [CORSRule](#) data types

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

CORSRule

Service: Amazon S3

Specifies a cross-origin access rule for an Amazon S3 bucket.

Contents

AllowedMethods

An HTTP method that you allow the origin to execute. Valid values are GET, PUT, HEAD, POST, and DELETE.

Type: Array of strings

Required: Yes

AllowedOrigins

One or more origins you want customers to be able to access the bucket from.

Type: Array of strings

Required: Yes

AllowedHeaders

Headers that are specified in the `Access-Control-Request-Headers` header. These headers are allowed in a preflight OPTIONS request. In response to any preflight OPTIONS request, Amazon S3 returns any requested headers that are allowed.

Type: Array of strings

Required: No

ExposeHeaders

One or more headers in the response that you want customers to be able to access from their applications (for example, from a JavaScript XMLHttpRequest object).

Type: Array of strings

Required: No

ID

Unique identifier for the rule. The value cannot be longer than 255 characters.

Type: String

Required: No

MaxAgeSeconds

The time in seconds that your browser is to cache the preflight response for the specified resource.

Type: Integer

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

CreateBucketConfiguration

Service: Amazon S3

The configuration information for the bucket.

Contents

Bucket

Specifies the information about the bucket that will be created.

Note

This functionality is only supported by directory buckets.

Type: [BucketInfo](#) data type

Required: No

Location

Specifies the location where the bucket will be created.

Directory buckets - The location type is Availability Zone or Local Zone. To use the Local Zone location type, your account must be enabled for Local Zones. Otherwise, you get an HTTP 403 Forbidden error with the error code AccessDenied. To learn more, see [Enable accounts for Local Zones](#) in the *Amazon S3 User Guide*.

Note

This functionality is only supported by directory buckets.

Type: [LocationInfo](#) data type

Required: No

LocationConstraint

Specifies the Region where the bucket will be created. You might choose a Region to optimize latency, minimize costs, or address regulatory requirements. For example, if you reside in Europe, you will probably find it advantageous to create buckets in the Europe (Ireland) Region.

If you don't specify a Region, the bucket is created in the US East (N. Virginia) Region (us-east-1) by default. Configurations using the value EU will create a bucket in eu-west-1.

For a list of the valid values for all of the AWS Regions, see [Regions and Endpoints](#).

 Note

This functionality is not supported for directory buckets.

Type: String

Valid Values: af-south-1 | ap-east-1 | ap-east-2 | ap-northeast-1 | ap-northeast-2 | ap-northeast-3 | ap-south-1 | ap-south-2 | ap-southeast-1 | ap-southeast-2 | ap-southeast-3 | ap-southeast-4 | ap-southeast-5 | ap-southeast-6 | ap-southeast-7 | ca-central-1 | ca-west-1 | cn-north-1 | cn-northwest-1 | EU | eu-central-1 | eu-central-2 | eu-north-1 | eu-south-1 | eu-south-2 | eu-west-1 | eu-west-2 | eu-west-3 | il-central-1 | me-central-1 | me-south-1 | mx-central-1 | sa-east-1 | us-east-2 | us-gov-east-1 | us-gov-west-1 | us-west-1 | us-west-2

Required: No

Tags

An array of tags that you can apply to the bucket that you're creating. Tags are key-value pairs of metadata used to categorize and organize your buckets, track costs, and control access.

You must have the `s3:TagResource` permission to create a general purpose bucket with tags or the `s3express:TagResource` permission to create a directory bucket with tags.

When creating buckets with tags, note that tag-based conditions using `aws:ResourceTag` and `s3:BucketTag` condition keys are applicable only after ABAC is enabled on the bucket. To learn more, see [Enabling ABAC in general purpose buckets](#).

Type: Array of [Tag](#) data types

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

CSVInput

Service: Amazon S3

Describes how an uncompressed comma-separated values (CSV)-formatted input object is formatted.

Contents

AllowQuotedRecordDelimiter

Specifies that CSV field values may contain quoted record delimiters and such records should be allowed. Default value is FALSE. Setting this value to TRUE may lower performance.

Type: Boolean

Required: No

Comments

A single character used to indicate that a row should be ignored when the character is present at the start of that row. You can specify any character to indicate a comment line. The default character is #.

Default: #

Type: String

Required: No

FieldDelimiter

A single character used to separate individual fields in a record. You can specify an arbitrary delimiter.

Type: String

Required: No

FileHeaderInfo

Describes the first line of input. Valid values are:

- NONE: First line is not a header.

- **IGNORE:** First line is a header, but you can't use the header values to indicate the column in an expression. You can use column position (such as `_1`, `_2`, ...) to indicate the column (`SELECT s._1 FROM OBJECT s`).
- **Use:** First line is a header, and you can use the header value to identify a column in an expression (`SELECT "name" FROM OBJECT`).

Type: String

Valid Values: `USE` | `IGNORE` | `NONE`

Required: No

QuoteCharacter

A single character used for escaping when the field delimiter is part of the value. For example, if the value is `a , b`, Amazon S3 wraps this field value in quotation marks, as follows: `" a , b "`.

Type: String

Default: `"`

Ancestors: `CSV`

Type: String

Required: No

QuoteEscapeCharacter

A single character used for escaping the quotation mark character inside an already escaped value. For example, the value `"" " a , b ""` is parsed as `" a , b "`.

Type: String

Required: No

RecordDelimiter

A single character used to separate individual records in the input. Instead of the default value, you can specify an arbitrary delimiter.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

CSVOutput

Service: Amazon S3

Describes how uncompressed comma-separated values (CSV)-formatted results are formatted.

Contents

FieldDelimiter

The value used to separate individual fields in a record. You can specify an arbitrary delimiter.

Type: String

Required: No

QuoteCharacter

A single character used for escaping when the field delimiter is part of the value. For example, if the value is a , b, Amazon S3 wraps this field value in quotation marks, as follows: " a , b ".

Type: String

Required: No

QuoteEscapeCharacter

The single character used for escaping the quote character inside an already escaped value.

Type: String

Required: No

QuoteFields

Indicates whether to use quotation marks around output fields.

- ALWAYS: Always use quotation marks for output fields.
- ASNEEDED: Use quotation marks for output fields when needed.

Type: String

Valid Values: ALWAYS | ASNEEDED

Required: No

RecordDelimiter

A single character used to separate individual records in the output. Instead of the default value, you can specify an arbitrary delimiter.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

DefaultRetention

Service: Amazon S3

The container element for optionally specifying the default Object Lock retention settings for new objects placed in the specified bucket.

Note

- The DefaultRetention settings require both a mode and a period.
- The DefaultRetention period can be either Days or Years but you must select one. You cannot specify Days and Years at the same time.

Contents

Days

The number of days that you want to specify for the default retention period. Must be used with Mode.

Type: Integer

Required: No

DefaultEventHold

The default event hold duration to be applied to new objects placed in the specified bucket. When configured, new objects will automatically have an event hold enabled with this duration.

Type: [EventHoldDuration](#) data type

Required: No

Mode

The default Object Lock retention mode you want to apply to new objects placed in the specified bucket. Must be used with either Days or Years.

Type: String

Valid Values: GOVERNANCE | COMPLIANCE

Required: No

Years

The number of years that you want to specify for the default retention period. Must be used with Mode.

Type: Integer

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Delete

Service: Amazon S3

Container for the objects to delete.

Contents

Objects

The object to delete.

Note

Directory buckets - For directory buckets, an object that's composed entirely of whitespace characters is not supported by the `DeleteObjects` API operation. The request will receive a `400 Bad Request` error and none of the objects in the request will be deleted.

Type: Array of [ObjectIdentifier](#) data types

Required: Yes

Quiet

Element to enable quiet mode for the request. When you add this element, you must set its value to `true`.

Type: Boolean

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

DeletedObject

Service: Amazon S3

Information about the deleted object.

Contents

DeleteMarker

Indicates whether the specified object version that was permanently deleted was (true) or was not (false) a delete marker before deletion. In a simple DELETE, this header indicates whether (true) or not (false) the current version of the object is a delete marker. To learn more about delete markers, see [Working with delete markers](#).

Note

This functionality is not supported for directory buckets.

Type: Boolean

Required: No

DeleteMarkerVersionId

The version ID of the delete marker created as a result of the DELETE operation. If you delete a specific object version, the value returned by this header is the version ID of the object version deleted.

Note

This functionality is not supported for directory buckets.

Type: String

Required: No

Key

The name of the deleted object.

Type: String

Length Constraints: Minimum length of 1.

Required: No

VersionId

The version ID of the deleted object.

Note

This functionality is not supported for directory buckets.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

DeleteMarkerEntry

Service: Amazon S3

Information about the delete marker.

Contents

IsLatest

Specifies whether the object is (true) or is not (false) the latest version of an object.

Type: Boolean

Required: No

Key

The object key.

Type: String

Length Constraints: Minimum length of 1.

Required: No

LastModified

Date and time when the object was last modified.

Type: Timestamp

Required: No

Owner

The account that created the delete marker.

Type: [Owner](#) data type

Required: No

VersionId

Version ID of an object.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

DeleteMarkerReplication

Service: Amazon S3

Specifies whether Amazon S3 replicates delete markers. If you specify a `Filter` in your replication configuration, you must also include a `DeleteMarkerReplication` element. If your `Filter` includes a `Tag` element, the `DeleteMarkerReplication Status` must be set to `Disabled`, because Amazon S3 does not support replicating delete markers for tag-based rules. For an example configuration, see [Basic Rule Configuration](#).

For more information about delete marker replication, see [Basic Rule Configuration](#).

Note

If you are using an earlier version of the replication configuration, Amazon S3 handles replication of delete markers differently. For more information, see [Backward Compatibility](#).

Contents

Status

Indicates whether to replicate delete markers.

Note

Indicates whether to replicate delete markers.

Type: String

Valid Values: `Enabled` | `Disabled`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Destination

Service: Amazon S3

Specifies information about where to publish analysis or configuration results for an Amazon S3 bucket and S3 Replication Time Control (S3 RTC).

Contents

Bucket

The Amazon Resource Name (ARN) of the bucket where you want Amazon S3 to store the results.

Type: String

Required: Yes

AccessControlTranslation

Specify this only in a cross-account scenario (where source and destination bucket owners are not the same), and you want to change replica ownership to the AWS account that owns the destination bucket. If this is not specified in the replication configuration, the replicas are owned by same AWS account that owns the source object.

Type: [AccessControlTranslation](#) data type

Required: No

Account

Destination bucket owner account ID. In a cross-account scenario, if you direct Amazon S3 to change replica ownership to the AWS account that owns the destination bucket by specifying the `AccessControlTranslation` property, this is the account ID of the destination bucket owner. For more information, see [Replication Additional Configuration: Changing the Replica Owner](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

EncryptionConfiguration

A container that provides information about encryption. If `SourceSelectionCriteria` is specified, you must specify this element.

Type: [EncryptionConfiguration](#) data type

Required: No

Metrics

A container specifying replication metrics-related settings enabling replication metrics and events.

Type: [Metrics](#) data type

Required: No

ReplicationTime

A container specifying S3 Replication Time Control (S3 RTC), including whether S3 RTC is enabled and the time when all objects and operations on objects must be replicated. Must be specified together with a `Metrics` block.

Type: [ReplicationTime](#) data type

Required: No

StorageClass

The storage class to use when replicating objects, such as S3 Standard or reduced redundancy. By default, Amazon S3 uses the storage class of the source object to create the object replica.

For valid values, see the `StorageClass` element of the [PUT Bucket replication](#) action in the *Amazon S3 API Reference*.

FSX_OPENZFS is not an accepted value when replicating objects.

Type: String

Valid Values: STANDARD | REDUCED_REDUNDANCY | STANDARD_IA | ONEZONE_IA | INTELLIGENT_TIERING | GLACIER | DEEP_ARCHIVE | OUTPOSTS | GLACIER_IR | SNOW | EXPRESS_ONEZONE | FSX_OPENZFS | FSX_ONTAP | AWS_BACKUP_WARM | AWS_BACKUP_LOW_COST_WARM

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

DestinationResult

Service: Amazon S3

The destination information for the S3 Metadata configuration.

Contents

TableBucketArn

The Amazon Resource Name (ARN) of the table bucket where the metadata configuration is stored.

Type: String

Required: No

TableBucketType

The type of the table bucket where the metadata configuration is stored. The `aws` value indicates an AWS managed table bucket, and the `customer` value indicates a customer-managed table bucket. V2 metadata configurations are stored in AWS managed table buckets, and V1 metadata configurations are stored in customer-managed table buckets.

Type: String

Valid Values: `aws` | `customer`

Required: No

TableNamespace

The namespace in the table bucket where the metadata tables for a metadata configuration are stored.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Encryption

Service: Amazon S3

Contains the type of server-side encryption used.

Contents

EncryptionType

The server-side encryption algorithm used when storing job results in Amazon S3 (for example, AES256, `aws:kms`).

Type: String

Valid Values: `AES256` | `aws:fsx` | `aws:backup` | `aws:kms` | `aws:kms:dsse`

Required: Yes

KMSContext

If the encryption type is `aws:kms`, this optional value can be used to specify the encryption context for the restore results.

Type: String

Required: No

KMSKeyId

If the encryption type is `aws:kms`, this optional value specifies the ID of the symmetric encryption customer managed key to use for encryption of job results. Amazon S3 only supports symmetric encryption KMS keys. For more information, see [Asymmetric keys in AWS KMS](#) in the *AWS Key Management Service Developer Guide*.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

EncryptionConfiguration

Service: Amazon S3

Specifies encryption-related information for an Amazon S3 bucket that is a destination for replicated objects.

Note

If you're specifying a customer managed KMS key, we recommend using a fully qualified KMS key ARN. If you use a KMS key alias instead, then AWS KMS resolves the key within the requester's account. This behavior can result in data that's encrypted with a KMS key that belongs to the requester, and not the bucket owner.

Contents

ReplicaKmsKeyID

Specifies the ID (Key ARN or Alias ARN) of the customer managed AWS KMS key stored in AWS Key Management Service (KMS) for the destination bucket. Amazon S3 uses this key to encrypt replica objects. Amazon S3 only supports symmetric encryption KMS keys. For more information, see [Asymmetric keys in AWS KMS](#) in the *AWS Key Management Service Developer Guide*.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

EndEvent

Service: Amazon S3

A message that indicates the request is complete and no more messages will be sent. You should not assume that the request is complete until the client receives an EndEvent.

Contents

The members of this exception structure are context-dependent.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Error

Service: Amazon S3

Note

For information about using the Amazon S3 API—including error handling—see the [Amazon S3 Developer Guide](#).

Container for all error elements.

Contents

Code

The error code is a string that uniquely identifies an error condition. It is meant to be read and understood by programs that detect and handle errors by type. The following is a list of Amazon S3 error codes. For more information, see [Error responses](#).

- • *Code:* AccessDenied
 - *Description:* Access Denied
 - *HTTP Status Code:* 403 Forbidden
 - *SOAP Fault Code Prefix:* Client
- • *Code:* AccountProblem
 - *Description:* There is a problem with your AWS account that prevents the action from completing successfully. Contact AWS Support for further assistance.
 - *HTTP Status Code:* 403 Forbidden
 - *SOAP Fault Code Prefix:* Client
- • *Code:* AllAccessDisabled
 - *Description:* All access to this Amazon S3 resource has been disabled. Contact AWS Support for further assistance.
 - *HTTP Status Code:* 403 Forbidden
 - *SOAP Fault Code Prefix:* Client
- • *Code:* AmbiguousGrantByEmailAddress
 - *Description:* The email address you provided is associated with more than one account.

- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* AuthorizationHeaderMalformed
 - *Description:* The authorization header you provided is invalid.
 - *HTTP Status Code:* 400 Bad Request
 - *HTTP Status Code:* N/A
- • *Code:* BadDigest
 - *Description:* The Content-MD5 you specified did not match what we received.
 - *HTTP Status Code:* 400 Bad Request
 - *SOAP Fault Code Prefix:* Client
- • *Code:* BucketAlreadyExists
 - *Description:* The requested bucket name is not available. The bucket namespace is shared by all users of the system. Please select a different name and try again.
 - *HTTP Status Code:* 409 Conflict
 - *SOAP Fault Code Prefix:* Client
- • *Code:* BucketAlreadyOwnedByYou
 - *Description:* The bucket you tried to create already exists, and you own it. Amazon S3 returns this error in all AWS Regions except in the North Virginia Region. For legacy compatibility, if you re-create an existing bucket that you already own in the North Virginia Region, Amazon S3 returns 200 OK and resets the bucket access control lists (ACLs).
 - *Code:* 409 Conflict (in all Regions except the North Virginia Region)
 - *SOAP Fault Code Prefix:* Client
- • *Code:* BucketNotEmpty
 - *Description:* The bucket you tried to delete is not empty.
 - *HTTP Status Code:* 409 Conflict
 - *SOAP Fault Code Prefix:* Client
- • *Code:* CredentialsNotSupported
 - *Description:* This request does not support credentials.
 - *HTTP Status Code:* 400 Bad Request
 - *SOAP Fault Code Prefix:* Client
- • *Code:* CrossLocationLoggingProhibited

- *Description:* Cross-location logging not allowed. Buckets in one geographic location cannot log information to a bucket in another location.
- *HTTP Status Code:* 403 Forbidden
- *SOAP Fault Code Prefix:* Client
- • *Code:* EntityTooSmall
- *Description:* Your proposed upload is smaller than the minimum allowed object size.
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* EntityTooLarge
- *Description:* Your proposed upload exceeds the maximum allowed object size.
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* ExpiredToken
- *Description:* The provided token has expired.
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* IllegalVersioningConfigurationException
- *Description:* Indicates that the versioning configuration specified in the request is invalid.
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* IncompleteBody
- *Description:* You did not provide the number of bytes specified by the Content-Length HTTP header
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* IncorrectNumberOfFilesInPostRequest
- *Description:* POST requires exactly one file upload per request.
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* InlineDataTooLarge
- *Description:* Inline data exceeds the maximum allowed size.

- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* InternalError
- *Description:* We encountered an internal error. Please try again.
- *HTTP Status Code:* 500 Internal Server Error
- *SOAP Fault Code Prefix:* Server
- • *Code:* InvalidAccessKeyId
- *Description:* The AWS access key ID you provided does not exist in our records.
- *HTTP Status Code:* 403 Forbidden
- *SOAP Fault Code Prefix:* Client
- • *Code:* InvalidAddressingHeader
- *Description:* You must specify the Anonymous role.
- *HTTP Status Code:* N/A
- *SOAP Fault Code Prefix:* Client
- • *Code:* InvalidArgument
- *Description:* Invalid Argument
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* InvalidBucketName
- *Description:* The specified bucket is not valid.
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* InvalidBucketState
- *Description:* The request is not valid with the current state of the bucket.
- *HTTP Status Code:* 409 Conflict
- *SOAP Fault Code Prefix:* Client
- • *Code:* InvalidDigest
- *Description:* The Content-MD5 you specified is not valid.
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client

- • *Code:* InvalidEncryptionAlgorithmError
 - *Description:* The encryption request you specified is not valid. The valid value is AES256.
 - *HTTP Status Code:* 400 Bad Request
 - *SOAP Fault Code Prefix:* Client
- • *Code:* InvalidLocationConstraint
 - *Description:* The specified location constraint is not valid. For more information about Regions, see [How to Select a Region for Your Buckets](#).
 - *HTTP Status Code:* 400 Bad Request
 - *SOAP Fault Code Prefix:* Client
- • *Code:* InvalidObjectState
 - *Description:* The action is not valid for the current state of the object.
 - *HTTP Status Code:* 403 Forbidden
 - *SOAP Fault Code Prefix:* Client
- • *Code:* InvalidPart
 - *Description:* One or more of the specified parts could not be found. The part might not have been uploaded, or the specified entity tag might not have matched the part's entity tag.
 - *HTTP Status Code:* 400 Bad Request
 - *SOAP Fault Code Prefix:* Client
- • *Code:* InvalidPartOrder
 - *Description:* The list of parts was not in ascending order. Parts list must be specified in order by part number.
 - *HTTP Status Code:* 400 Bad Request
 - *SOAP Fault Code Prefix:* Client
- • *Code:* InvalidPayer
 - *Description:* All access to this object has been disabled. Please contact AWS Support for further assistance.
 - *HTTP Status Code:* 403 Forbidden
 - *SOAP Fault Code Prefix:* Client
- • *Code:* InvalidPolicyDocument

- *Description:* The content of the form does not meet the conditions specified in the policy document.
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* InvalidRange
- *Description:* The requested range cannot be satisfied.
- *HTTP Status Code:* 416 Requested Range Not Satisfiable
- *SOAP Fault Code Prefix:* Client
- • *Code:* InvalidRequest
- *Description:* Please use AWS4-HMAC-SHA256.
- *HTTP Status Code:* 400 Bad Request
- *Code:* N/A
- • *Code:* InvalidRequest
- *Description:* SOAP requests must be made over an HTTPS connection.
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* InvalidRequest
- *Description:* Amazon S3 Transfer Acceleration is not supported for buckets with non-DNS compliant names.
- *HTTP Status Code:* 400 Bad Request
- *Code:* N/A
- • *Code:* InvalidRequest
- *Description:* Amazon S3 Transfer Acceleration is not supported for buckets with periods (.) in their names.
- *HTTP Status Code:* 400 Bad Request
- *Code:* N/A
- • *Code:* InvalidRequest
- *Description:* Amazon S3 Transfer Accelerate endpoint only supports virtual style requests.
- *HTTP Status Code:* 400 Bad Request
- *Code:* N/A
- • *Code:* InvalidRequest

- *Description:* Amazon S3 Transfer Accelerate is not configured on this bucket.
- *HTTP Status Code:* 400 Bad Request
- *Code:* N/A
- • *Code:* InvalidRequest
- *Description:* Amazon S3 Transfer Accelerate is disabled on this bucket.
- *HTTP Status Code:* 400 Bad Request
- *Code:* N/A
- • *Code:* InvalidRequest
- *Description:* Amazon S3 Transfer Acceleration is not supported on this bucket. Contact AWS Support for more information.
- *HTTP Status Code:* 400 Bad Request
- *Code:* N/A
- • *Code:* InvalidRequest
- *Description:* Amazon S3 Transfer Acceleration cannot be enabled on this bucket. Contact AWS Support for more information.
- *HTTP Status Code:* 400 Bad Request
- *Code:* N/A
- • *Code:* InvalidSecurity
- *Description:* The provided security credentials are not valid.
- *HTTP Status Code:* 403 Forbidden
- *SOAP Fault Code Prefix:* Client
- • *Code:* InvalidSOAPRequest
- *Description:* The SOAP request body is invalid.
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* InvalidStorageClass
- *Description:* The storage class you specified is not valid.
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client

- *Description:* The target bucket for logging does not exist, is not owned by you, or does not have the appropriate grants for the log-delivery group.
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* InvalidToken
- *Description:* The provided token is malformed or otherwise invalid.
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* InvalidURI
- *Description:* Couldn't parse the specified URI.
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* KeyTooLongError
- *Description:* Your key is too long.
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* MalformedACLError
- *Description:* The XML you provided was not well-formed or did not validate against our published schema.
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* MalformedPOSTRequest
- *Description:* The body of your POST request is not well-formed multipart/form-data.
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* MalformedXML
- *Description:* This happens when the user sends malformed XML (XML that doesn't conform to the published XSD) for the configuration. The error message is, "The XML you provided was not well-formed or did not validate against our published schema."
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client

- • *Code:* MaxMessageLengthExceeded
 - *Description:* Your request was too big.
 - *HTTP Status Code:* 400 Bad Request
 - *SOAP Fault Code Prefix:* Client
- • *Code:* MaxPostPreDataLengthExceededError
 - *Description:* Your POST request fields preceding the upload file were too large.
 - *HTTP Status Code:* 400 Bad Request
 - *SOAP Fault Code Prefix:* Client
- • *Code:* MetadataTooLarge
 - *Description:* Your metadata headers exceed the maximum allowed metadata size.
 - *HTTP Status Code:* 400 Bad Request
 - *SOAP Fault Code Prefix:* Client
- • *Code:* MethodNotAllowed
 - *Description:* The specified method is not allowed against this resource.
 - *HTTP Status Code:* 405 Method Not Allowed
 - *SOAP Fault Code Prefix:* Client
- • *Code:* MissingAttachment
 - *Description:* A SOAP attachment was expected, but none were found.
 - *HTTP Status Code:* N/A
 - *SOAP Fault Code Prefix:* Client
- • *Code:* MissingContentLength
 - *Description:* You must provide the Content-Length HTTP header.
 - *HTTP Status Code:* 411 Length Required
 - *SOAP Fault Code Prefix:* Client
- • *Code:* MissingRequestBodyError
 - *Description:* This happens when the user sends an empty XML document as a request. The error message is, "Request body is empty."
 - *HTTP Status Code:* 400 Bad Request
 - *SOAP Fault Code Prefix:* Client
- • *Code:* MissingSecurityElement
 - *Description:* The SOAP 1.1 request is missing a security element.

- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* MissingSecurityHeader
- *Description:* Your request is missing a required header.
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* NoLoggingStatusForKey
- *Description:* There is no such thing as a logging status subresource for a key.
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* NoSuchBucket
- *Description:* The specified bucket does not exist.
- *HTTP Status Code:* 404 Not Found
- *SOAP Fault Code Prefix:* Client
- • *Code:* NoSuchBucketPolicy
- *Description:* The specified bucket does not have a bucket policy.
- *HTTP Status Code:* 404 Not Found
- *SOAP Fault Code Prefix:* Client
- • *Code:* NoSuchKey
- *Description:* The specified key does not exist.
- *HTTP Status Code:* 404 Not Found
- *SOAP Fault Code Prefix:* Client
- • *Code:* NoSuchLifecycleConfiguration
- *Description:* The lifecycle configuration does not exist.
- *HTTP Status Code:* 404 Not Found
- *SOAP Fault Code Prefix:* Client
- • *Code:* NoSuchUpload
- *Description:* The specified multipart upload does not exist. The upload ID might be invalid, or the multipart upload might have been aborted or completed.
- *HTTP Status Code:* 404 Not Found

- *SOAP Fault Code Prefix:* Client
- • *Code:* NoSuchVersion
 - *Description:* Indicates that the version ID specified in the request does not match an existing version.
 - *HTTP Status Code:* 404 Not Found
 - *SOAP Fault Code Prefix:* Client
- • *Code:* NotImplemented
 - *Description:* A header you provided implies functionality that is not implemented.
 - *HTTP Status Code:* 501 Not Implemented
 - *SOAP Fault Code Prefix:* Server
- • *Code:* NotSignedUp
 - *Description:* Your account is not signed up for the Amazon S3 service. You must sign up before you can use Amazon S3. You can sign up at the following URL: [Amazon S3](#)
 - *HTTP Status Code:* 403 Forbidden
 - *SOAP Fault Code Prefix:* Client
- • *Code:* OperationAborted
 - *Description:* A conflicting conditional action is currently in progress against this resource. Try again.
 - *HTTP Status Code:* 409 Conflict
 - *SOAP Fault Code Prefix:* Client
- • *Code:* PermanentRedirect
 - *Description:* The bucket you are attempting to access must be addressed using the specified endpoint. Send all future requests to this endpoint.
 - *HTTP Status Code:* 301 Moved Permanently
 - *SOAP Fault Code Prefix:* Client
- • *Code:* PreconditionFailed
 - *Description:* At least one of the preconditions you specified did not hold.
 - *HTTP Status Code:* 412 Precondition Failed
 - *SOAP Fault Code Prefix:* Client
- • *Code:* Redirect
 - *Description:* Temporary redirect.

- *HTTP Status Code:* 307 Moved Temporarily
- *SOAP Fault Code Prefix:* Client
- • *Code:* RestoreAlreadyInProgress
- *Description:* Object restore is already in progress.
- *HTTP Status Code:* 409 Conflict
- *SOAP Fault Code Prefix:* Client
- • *Code:* RequestIsNotMultiPartContent
- *Description:* Bucket POST must be of the enclosure-type multipart/form-data.
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* RequestTimeout
- *Description:* Your socket connection to the server was not read from or written to within the timeout period.
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* RequestTimeTooSkewed
- *Description:* The difference between the request time and the server's time is too large.
- *HTTP Status Code:* 403 Forbidden
- *SOAP Fault Code Prefix:* Client
- • *Code:* RequestTorrentOfBucketError
- *Description:* Requesting the torrent file of a bucket is not permitted.
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* SignatureDoesNotMatch
- *Description:* The request signature we calculated does not match the signature you provided. Check your AWS secret access key and signing method. For more information, see [REST Authentication](#) and [SOAP Authentication](#) for details.
- *HTTP Status Code:* 403 Forbidden
- *SOAP Fault Code Prefix:* Client
- • *Code:* ServiceUnavailable
- *Description:* Service is unable to handle request.

- *HTTP Status Code:* 503 Service Unavailable
- *SOAP Fault Code Prefix:* Server
- • *Code:* SlowDown
- *Description:* Reduce your request rate.
- *HTTP Status Code:* 503 Slow Down
- *SOAP Fault Code Prefix:* Server
- • *Code:* TemporaryRedirect
- *Description:* You are being redirected to the bucket while DNS updates.
- *HTTP Status Code:* 307 Moved Temporarily
- *SOAP Fault Code Prefix:* Client
- • *Code:* TokenRefreshRequired
- *Description:* The provided token must be refreshed.
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* TooManyBuckets
- *Description:* You have attempted to create more buckets than allowed.
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* UnexpectedContent
- *Description:* This request does not support content.
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* UnresolvableGrantByEmailAddress
- *Description:* The email address you provided does not match any account on record.
- *HTTP Status Code:* 400 Bad Request
- *SOAP Fault Code Prefix:* Client
- • *Code:* UserKeyMustBeSpecified
- *Description:* The bucket POST must contain the specified field name. If it is specified, check the order of the fields.
- *HTTP Status Code:* 400 Bad Request

- *SOAP Fault Code Prefix: Client*

Type: String

Required: No

Key

The error key.

Type: String

Length Constraints: Minimum length of 1.

Required: No

Message

The error message contains a generic description of the error condition in English. It is intended for a human audience. Simple programs display the message directly to the end user if they encounter an error condition they don't know how or don't care to handle. Sophisticated programs with more exhaustive error handling and proper internationalization are more likely to ignore the error message.

Type: String

Required: No

VersionId

The version ID of the error.

Note

This functionality is not supported for directory buckets.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ErrorDetails

Service: Amazon S3

If an S3 Metadata V1 `CreateBucketMetadataTableConfiguration` or V2 `CreateBucketMetadataConfiguration` request succeeds, but S3 Metadata was unable to create the table, this structure contains the error code and error message.

Note

If you created your S3 Metadata configuration before July 15, 2025, we recommend that you delete and re-create your configuration by using [CreateBucketMetadataConfiguration](#) so that you can expire journal table records and create a live inventory table.

Contents

ErrorCode

If the V1 `CreateBucketMetadataTableConfiguration` request succeeds, but S3 Metadata was unable to create the table, this structure contains the error code. The possible error codes and error messages are as follows:

- `AccessDeniedCreatingResources` - You don't have sufficient permissions to create the required resources. Make sure that you have `s3tables:CreateNamespace`, `s3tables:CreateTable`, `s3tables:GetTable` and `s3tables:PutTablePolicy` permissions, and then try again. To create a new metadata table, you must delete the metadata configuration for this bucket, and then create a new metadata configuration.
- `AccessDeniedWritingToTable` - Unable to write to the metadata table because of missing resource permissions. To fix the resource policy, Amazon S3 needs to create a new metadata table. To create a new metadata table, you must delete the metadata configuration for this bucket, and then create a new metadata configuration.
- `DestinationTableNotFound` - The destination table doesn't exist. To create a new metadata table, you must delete the metadata configuration for this bucket, and then create a new metadata configuration.
- `ServerInternalError` - An internal error has occurred. To create a new metadata table, you must delete the metadata configuration for this bucket, and then create a new metadata configuration.

- **TableAlreadyExists** - The table that you specified already exists in the table bucket's namespace. Specify a different table name. To create a new metadata table, you must delete the metadata configuration for this bucket, and then create a new metadata configuration.
- **TableBucketNotFound** - The table bucket that you specified doesn't exist in this AWS Region and account. Create or choose a different table bucket. To create a new metadata table, you must delete the metadata configuration for this bucket, and then create a new metadata configuration.

If the V2 `CreateBucketMetadataConfiguration` request succeeds, but S3 Metadata was unable to create the table, this structure contains the error code. The possible error codes and error messages are as follows:

- **AccessDeniedCreatingResources** - You don't have sufficient permissions to create the required resources. Make sure that you have `s3tables:CreateTableBucket`, `s3tables:CreateNamespace`, `s3tables:CreateTable`, `s3tables:GetTable`, `s3tables:PutTablePolicy`, `kms:DescribeKey`, and `s3tables:PutTableEncryption` permissions. Additionally, ensure that the KMS key used to encrypt the table still exists, is active and has a resource policy granting access to the S3 service principals `'maintenance.s3tables.amazonaws.com'` and `'metadata.s3.amazonaws.com'`. To create a new metadata table, you must delete the metadata configuration for this bucket, and then create a new metadata configuration.
- **AccessDeniedWritingToTable** - Unable to write to the metadata table because of missing resource permissions. To fix the resource policy, Amazon S3 needs to create a new metadata table. To create a new metadata table, you must delete the metadata configuration for this bucket, and then create a new metadata configuration.
- **DestinationTableNotFound** - The destination table doesn't exist. To create a new metadata table, you must delete the metadata configuration for this bucket, and then create a new metadata configuration.
- **ServerInternalError** - An internal error has occurred. To create a new metadata table, you must delete the metadata configuration for this bucket, and then create a new metadata configuration.
- **JournalTableAlreadyExists** - A journal table already exists in the AWS managed table bucket's namespace. Delete the journal table, and then try again. To create a new metadata table, you must delete the metadata configuration for this bucket, and then create a new metadata configuration.

- `InventoryTableAlreadyExists` - An inventory table already exists in the AWS managed table bucket's namespace. Delete the inventory table, and then try again. To create a new metadata table, you must delete the metadata configuration for this bucket, and then create a new metadata configuration.
- `JournalTableNotAvailable` - The journal table that the inventory table relies on has a `FAILED` status. An inventory table requires a journal table with an `ACTIVE` status. To create a new journal or inventory table, you must delete the metadata configuration for this bucket, along with any journal or inventory tables, and then create a new metadata configuration.
- `NoSuchBucket` - The specified general purpose bucket does not exist.

Type: String

Required: No

Error Message

If the V1 `CreateBucketMetadataTableConfiguration` request succeeds, but S3 Metadata was unable to create the table, this structure contains the error message. The possible error codes and error messages are as follows:

- `AccessDeniedCreatingResources` - You don't have sufficient permissions to create the required resources. Make sure that you have `s3tables:CreateNamespace`, `s3tables:CreateTable`, `s3tables:GetTable` and `s3tables:PutTablePolicy` permissions, and then try again. To create a new metadata table, you must delete the metadata configuration for this bucket, and then create a new metadata configuration.
- `AccessDeniedWritingToTable` - Unable to write to the metadata table because of missing resource permissions. To fix the resource policy, Amazon S3 needs to create a new metadata table. To create a new metadata table, you must delete the metadata configuration for this bucket, and then create a new metadata configuration.
- `DestinationTableNotFound` - The destination table doesn't exist. To create a new metadata table, you must delete the metadata configuration for this bucket, and then create a new metadata configuration.
- `ServerInternalError` - An internal error has occurred. To create a new metadata table, you must delete the metadata configuration for this bucket, and then create a new metadata configuration.
- `TableAlreadyExists` - The table that you specified already exists in the table bucket's namespace. Specify a different table name. To create a new metadata table, you must delete the metadata configuration for this bucket, and then create a new metadata configuration.

- **TableBucketNotFound** - The table bucket that you specified doesn't exist in this AWS Region and account. Create or choose a different table bucket. To create a new metadata table, you must delete the metadata configuration for this bucket, and then create a new metadata configuration.

If the V2 `CreateBucketMetadataConfiguration` request succeeds, but S3 Metadata was unable to create the table, this structure contains the error code. The possible error codes and error messages are as follows:

- **AccessDeniedCreatingResources** - You don't have sufficient permissions to create the required resources. Make sure that you have `s3tables:CreateTableBucket`, `s3tables:CreateNamespace`, `s3tables:CreateTable`, `s3tables:GetTable`, `s3tables:PutTablePolicy`, `kms:DescribeKey`, and `s3tables:PutTableEncryption` permissions. Additionally, ensure that the KMS key used to encrypt the table still exists, is active and has a resource policy granting access to the S3 service principals `'maintenance.s3tables.amazonaws.com'` and `'metadata.s3.amazonaws.com'`. To create a new metadata table, you must delete the metadata configuration for this bucket, and then create a new metadata configuration.
- **AccessDeniedWritingToTable** - Unable to write to the metadata table because of missing resource permissions. To fix the resource policy, Amazon S3 needs to create a new metadata table. To create a new metadata table, you must delete the metadata configuration for this bucket, and then create a new metadata configuration.
- **DestinationTableNotFound** - The destination table doesn't exist. To create a new metadata table, you must delete the metadata configuration for this bucket, and then create a new metadata configuration.
- **ServerInternalError** - An internal error has occurred. To create a new metadata table, you must delete the metadata configuration for this bucket, and then create a new metadata configuration.
- **JournalTableAlreadyExists** - A journal table already exists in the AWS managed table bucket's namespace. Delete the journal table, and then try again. To create a new metadata table, you must delete the metadata configuration for this bucket, and then create a new metadata configuration.
- **InventoryTableAlreadyExists** - An inventory table already exists in the AWS managed table bucket's namespace. Delete the inventory table, and then try again. To create a new metadata table, you must delete the metadata configuration for this bucket, and then create a new metadata configuration.

- `JournalTableNotAvailable` - The journal table that the inventory table relies on has a `FAILED` status. An inventory table requires a journal table with an `ACTIVE` status. To create a new journal or inventory table, you must delete the metadata configuration for this bucket, along with any journal or inventory tables, and then create a new metadata configuration.
- `NoSuchBucket` - The specified general purpose bucket does not exist.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ErrorDocument

Service: Amazon S3

The error information.

Contents

Key

The object key name to use when a 4XX class error occurs.

Important

Replacement must be made for object keys containing special characters (such as carriage returns) when using XML requests. For more information, see [XML related object key constraints](#).

Type: String

Length Constraints: Minimum length of 1.

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

EventBridgeConfiguration

Service: Amazon S3

A container for specifying the configuration for Amazon EventBridge.

Contents

The members of this exception structure are context-dependent.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

EventHoldDuration

Service: Amazon S3

Contains the event hold duration configuration, specified in either days or years.

Contents

Days

The number of days for the event hold duration. The minimum value is 1 and the maximum value is 36,500.

Type: Integer

Required: No

Years

The number of years for the event hold duration. The minimum value is 1 and the maximum value is 100.

Type: Integer

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ExistingObjectReplication

Service: Amazon S3

Optional configuration to replicate existing source bucket objects.

Note

This parameter is no longer supported. To replicate existing objects, see [Replicating existing objects with S3 Batch Replication](#) in the *Amazon S3 User Guide*.

Contents

Status

Specifies whether Amazon S3 replicates existing source bucket objects.

Type: String

Valid Values: Enabled | Disabled

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

FilterRule

Service: Amazon S3

Specifies the Amazon S3 object key name to filter on. An object key name is the name assigned to an object in your Amazon S3 bucket. You specify whether to filter on the suffix or prefix of the object key name. A prefix is a specific string of characters at the beginning of an object key name, which you can use to organize objects. For example, you can start the key names of related objects with a prefix, such as 2023- or engineering/. Then, you can use `FilterRule` to find objects in a bucket with key names that have the same prefix. A suffix is similar to a prefix, but it is at the end of the object key name instead of at the beginning.

Contents

Name

The object key name prefix or suffix identifying one or more objects to which the filtering rule applies. The maximum length is 1,024 characters. Overlapping prefixes and suffixes are not supported. For more information, see [Configuring Event Notifications](#) in the *Amazon S3 User Guide*.

Type: String

Valid Values: `prefix` | `suffix`

Required: No

Value

The value that the filter searches for in object key names.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)

- [AWS SDK for Ruby V3](#)

GetBucketMetadataConfigurationResult

Service: Amazon S3

The S3 Metadata configuration for a general purpose bucket.

Contents

MetadataConfigurationResult

The metadata configuration for a general purpose bucket.

Type: [MetadataConfigurationResult](#) data type

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

GetBucketMetadataTableConfigurationResult

Service: Amazon S3

The V1 S3 Metadata configuration for a general purpose bucket.

Note

If you created your S3 Metadata configuration before July 15, 2025, we recommend that you delete and re-create your configuration by using [CreateBucketMetadataConfiguration](#) so that you can expire journal table records and create a live inventory table.

Contents

MetadataTableConfigurationResult

The V1 S3 Metadata configuration for a general purpose bucket.

Type: [MetadataTableConfigurationResult](#) data type

Required: Yes

Status

The status of the metadata table. The status values are:

- **CREATING** - The metadata table is in the process of being created in the specified table bucket.
- **ACTIVE** - The metadata table has been created successfully, and records are being delivered to the table.
- **FAILED** - Amazon S3 is unable to create the metadata table, or Amazon S3 is unable to deliver records. See `ErrorDetails` for details.

Type: String

Required: Yes

Error

If the `CreateBucketMetadataTableConfiguration` request succeeds, but S3 Metadata was unable to create the table, this structure contains the error code and error message.

Type: [ErrorDetails](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

GetObjectAttributesParts

Service: Amazon S3

A collection of parts associated with a multipart upload.

Contents

IsTruncated

Indicates whether the returned list of parts is truncated. A value of `true` indicates that the list was truncated. A list can be truncated if the number of parts exceeds the limit returned in the `MaxParts` element.

Type: Boolean

Required: No

MaxParts

The maximum number of parts allowed in the response.

Type: Integer

Required: No

NextPartNumberMarker

When a list is truncated, this element specifies the last part in the list, as well as the value to use for the `PartNumberMarker` request parameter in a subsequent request.

Type: Integer

Required: No

PartNumberMarker

The marker for the current part.

Type: Integer

Required: No

Parts

A container for elements related to a particular part. A response can contain zero or more `Parts` elements.

Note

- **General purpose buckets** - For `GetObjectAttributes`, if an additional checksum (including `x-amz-checksum-crc32`, `x-amz-checksum-crc32c`, `x-amz-checksum-sha1`, or `x-amz-checksum-sha256`) isn't applied to the object specified in the request, the response doesn't return the `Part` element.
- **Directory buckets** - For `GetObjectAttributes`, regardless of whether an additional checksum is applied to the object specified in the request, the response returns the `Part` element.

Type: Array of [ObjectPart](#) data types

Required: No

TotalPartsCount

The total number of parts.

Type: Integer

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

GlacierJobParameters

Service: Amazon S3

Container for S3 Glacier job parameters.

Contents

Tier

Retrieval tier at which the restore will be processed.

Type: String

Valid Values: Standard | Bulk | Expedited

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Grant

Service: Amazon S3

Container for grant information.

Contents

Grantee

The person being granted permissions.

Type: [Grantee](#) data type

Required: No

Permission

Specifies the permission given to the grantee.

Type: String

Valid Values: FULL_CONTROL | WRITE | WRITE_ACP | READ | READ_ACP

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Grantee

Service: Amazon S3

Container for the person being granted permissions.

Contents

Type

Type of grantee

Type: String

Valid Values: CanonicalUser | AmazonCustomerByEmail | Group

Required: Yes

DisplayName

Type: String

Required: No

EmailAddress

Type: String

Required: No

ID

The canonical user ID of the grantee.

Type: String

Required: No

URI

URI of the grantee group.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

IndexDocument

Service: Amazon S3

Container for the `Suffix` element.

Contents

Suffix

A suffix that is appended to a request that is for a directory on the website endpoint. (For example, if the suffix is `index.html` and you make a request to `samplebucket/images/`, the data that is returned will be for the object with the key name `images/index.html`.) The suffix must not be empty and must not include a slash character.

Important

Replacement must be made for object keys containing special characters (such as carriage returns) when using XML requests. For more information, see [XML related object key constraints](#).

Type: String

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Initiator

Service: Amazon S3

Container element that identifies who initiated the multipart upload.

Contents

DisplayName

Note

This functionality is not supported for directory buckets.

Type: String

Required: No

ID

If the principal is an AWS account, it provides the Canonical User ID. If the principal is an IAM User, it provides a user ARN value.

Note

Directory buckets - If the principal is an AWS account, it provides the AWS account ID. If the principal is an IAM User, it provides a user ARN value.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

InputSerialization

Service: Amazon S3

Describes the serialization format of the object.

Contents

CompressionType

Specifies object's compression format. Valid values: NONE, GZIP, BZIP2. Default Value: NONE.

Type: String

Valid Values: NONE | GZIP | BZIP2

Required: No

CSV

Describes the serialization of a CSV-encoded object.

Type: [CSVInput](#) data type

Required: No

JSON

Specifies JSON as object's input serialization format.

Type: [JSONInput](#) data type

Required: No

Parquet

Specifies Parquet as object's input serialization format.

Type: [ParquetInput](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

IntelligentTieringAndOperator

Service: Amazon S3

A container for specifying S3 Intelligent-Tiering filters. The filters determine the subset of objects to which the rule applies.

Contents

Prefix

An object key name prefix that identifies the subset of objects to which the configuration applies.

Type: String

Required: No

Tags

All of these tags must exist in the object's tag set in order for the configuration to apply.

Type: Array of [Tag](#) data types

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

IntelligentTieringConfiguration

Service: Amazon S3

Specifies the S3 Intelligent-Tiering configuration for an Amazon S3 bucket.

For information about the S3 Intelligent-Tiering storage class, see [Storage class for automatically optimizing frequently and infrequently accessed objects](#).

Contents

Id

The ID used to identify the S3 Intelligent-Tiering configuration.

Type: String

Required: Yes

Status

Specifies the status of the configuration.

Type: String

Valid Values: Enabled | Disabled

Required: Yes

Tierings

Specifies the S3 Intelligent-Tiering storage class tier of the configuration.

Type: Array of [Tiering](#) data types

Required: Yes

Filter

Specifies a bucket filter. The configuration only includes objects that meet the filter's criteria.

Type: [IntelligentTieringFilter](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

IntelligentTieringFilter

Service: Amazon S3

The `Filter` is used to identify objects that the S3 Intelligent-Tiering configuration applies to.

Contents

And

A conjunction (logical AND) of predicates, which is used in evaluating a metrics filter. The operator must have at least two predicates, and an object must match all of the predicates in order for the filter to apply.

Type: [IntelligentTieringAndOperator](#) data type

Required: No

Prefix

An object key name prefix that identifies the subset of objects to which the rule applies.

Important

Replacement must be made for object keys containing special characters (such as carriage returns) when using XML requests. For more information, see [XML related object key constraints](#).

Type: String

Required: No

Tag

A container of a key value name pair.

Type: [Tag](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

InventoryConfiguration

Service: Amazon S3

Specifies the S3 Inventory configuration for an Amazon S3 bucket. For more information, see [GET Bucket inventory](#) in the *Amazon S3 API Reference*.

Contents

Destination

Contains information about where to publish the inventory results.

Type: [InventoryDestination](#) data type

Required: Yes

Id

The ID used to identify the inventory configuration.

Type: String

Required: Yes

IncludedObjectVersions

Object versions to include in the inventory list. If set to `All`, the list includes all the object versions, which adds the version-related fields `VersionId`, `IsLatest`, and `DeleteMarker` to the list. If set to `Current`, the list does not contain these version-related fields.

Type: String

Valid Values: `All` | `Current`

Required: Yes

IsEnabled

Specifies whether the inventory is enabled or disabled. If set to `True`, an inventory list is generated. If set to `False`, no inventory list is generated.

Type: Boolean

Required: Yes

Schedule

Specifies the schedule for generating inventory results.

Type: [InventorySchedule](#) data type

Required: Yes

Filter

Specifies an inventory filter. The inventory only includes objects that meet the filter's criteria.

Type: [InventoryFilter](#) data type

Required: No

OptionalFields

Contains the optional fields that are included in the inventory results.

Note

The following optional fields are supported for directory buckets Size | LastModifiedDate | StorageClass | ETag | IsMultipartUploaded | EncryptionStatus | BucketKeyStatus | ChecksumAlgorithm | LifecycleExpirationDate. Throws MalformedXML error if unsupported optional field is provided.

Type: Array of strings

Valid Values: Size | LastModifiedDate | StorageClass | ETag | IsMultipartUploaded | ReplicationStatus | EncryptionStatus | ObjectLockRetainUntilDate | ObjectLockMode | ObjectLockLegalHoldStatus | ObjectLockEventHoldStatus | ObjectLockEventHoldDuration | IntelligentTieringAccessTier | BucketKeyStatus | ChecksumAlgorithm | ObjectAccessControlList | ObjectOwner | LifecycleExpirationDate

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

InventoryDestination

Service: Amazon S3

Specifies the S3 Inventory configuration for an Amazon S3 bucket.

Contents

S3BucketDestination

Contains the bucket name, file format, bucket owner (optional), and prefix (optional) where inventory results are published.

Type: [InventoryS3BucketDestination](#) data type

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

InventoryEncryption

Service: Amazon S3

Contains the type of server-side encryption used to encrypt the S3 Inventory results.

Contents

SSEKMS

Specifies the use of SSE-KMS to encrypt delivered inventory reports.

Type: [SSEKMS](#) data type

Required: No

SSES3

Specifies the use of SSE-S3 to encrypt delivered inventory reports.

Type: [SSES3](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

InventoryFilter

Service: Amazon S3

Specifies an S3 Inventory filter. The inventory only includes objects that meet the filter's criteria.

Contents

Prefix

The prefix that an object must have to be included in the inventory results.

Type: String

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

InventoryS3BucketDestination

Service: Amazon S3

Contains the bucket name, file format, bucket owner (optional), and prefix (optional) where S3 Inventory results are published.

Contents

Bucket

The Amazon Resource Name (ARN) of the bucket where inventory results will be published.

Type: String

Required: Yes

Format

Specifies the output format of the inventory results.

Type: String

Valid Values: CSV | ORC | Parquet

Required: Yes

AccountId

The account ID that owns the destination S3 bucket. If no account ID is provided, the owner is not validated before exporting data.

Note

Although this value is optional, we strongly recommend that you set it to help prevent problems if the destination bucket ownership changes.

Type: String

Required: No

Encryption

Contains the type of server-side encryption used to encrypt the inventory results.

Type: [InventoryEncryption](#) data type

Required: No

Prefix

The prefix that is prepended to all inventory results.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

InventorySchedule

Service: Amazon S3

Specifies the schedule for generating S3 Inventory results.

Contents

Frequency

Specifies how frequently inventory results are produced.

Type: String

Valid Values: `Daily` | `Weekly`

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

InventoryTableConfiguration

Service: Amazon S3

The inventory table configuration for an S3 Metadata configuration.

Contents

ConfigurationState

The configuration state of the inventory table, indicating whether the inventory table is enabled or disabled.

Type: String

Valid Values: ENABLED | DISABLED

Required: Yes

EncryptionConfiguration

The encryption configuration for the inventory table.

Type: [MetadataTableEncryptionConfiguration](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

InventoryTableConfigurationResult

Service: Amazon S3

The inventory table configuration for an S3 Metadata configuration.

Contents

ConfigurationState

The configuration state of the inventory table, indicating whether the inventory table is enabled or disabled.

Type: String

Valid Values: ENABLED | DISABLED

Required: Yes

Error

If an S3 Metadata V1 `CreateBucketMetadataTableConfiguration` or V2 `CreateBucketMetadataConfiguration` request succeeds, but S3 Metadata was unable to create the table, this structure contains the error code and error message.

Note

If you created your S3 Metadata configuration before July 15, 2025, we recommend that you delete and re-create your configuration by using [CreateBucketMetadataConfiguration](#) so that you can expire journal table records and create a live inventory table.

Type: [ErrorDetails](#) data type

Required: No

TableArn

The Amazon Resource Name (ARN) for the inventory table.

Type: String

Required: No

TableName

The name of the inventory table.

Type: String

Required: No

TableStatus

The status of the inventory table. The status values are:

- **CREATING** - The inventory table is in the process of being created in the specified AWS managed table bucket.
- **BACKFILLING** - The inventory table is in the process of being backfilled. When you enable the inventory table for your metadata configuration, the table goes through a process known as backfilling, during which Amazon S3 scans your general purpose bucket to retrieve the initial metadata for all objects in the bucket. Depending on the number of objects in your bucket, this process can take several hours. When the backfilling process is finished, the status of your inventory table changes from **BACKFILLING** to **ACTIVE**. After backfilling is completed, updates to your objects are reflected in the inventory table within one hour.
- **ACTIVE** - The inventory table has been created successfully, and records are being delivered to the table.
- **FAILED** - Amazon S3 is unable to create the inventory table, or Amazon S3 is unable to deliver records.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

InventoryTableConfigurationUpdates

Service: Amazon S3

The specified updates to the S3 Metadata inventory table configuration.

Contents

ConfigurationState

The configuration state of the inventory table, indicating whether the inventory table is enabled or disabled.

Type: String

Valid Values: ENABLED | DISABLED

Required: Yes

EncryptionConfiguration

The encryption configuration for the inventory table.

Type: [MetadataTableEncryptionConfiguration](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

JournalTableConfiguration

Service: Amazon S3

The journal table configuration for an S3 Metadata configuration.

Contents

RecordExpiration

The journal table record expiration settings for the journal table.

Type: [RecordExpiration](#) data type

Required: Yes

EncryptionConfiguration

The encryption configuration for the journal table.

Type: [MetadataTableEncryptionConfiguration](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

JournalTableConfigurationResult

Service: Amazon S3

The journal table configuration for the S3 Metadata configuration.

Contents

RecordExpiration

The journal table record expiration settings for the journal table.

Type: [RecordExpiration](#) data type

Required: Yes

TableName

The name of the journal table.

Type: String

Required: Yes

TableStatus

The status of the journal table. The status values are:

- **CREATING** - The journal table is in the process of being created in the specified table bucket.
- **ACTIVE** - The journal table has been created successfully, and records are being delivered to the table.
- **FAILED** - Amazon S3 is unable to create the journal table, or Amazon S3 is unable to deliver records.

Type: String

Required: Yes

Error

If an S3 Metadata V1 `CreateBucketMetadataTableConfiguration` or V2 `CreateBucketMetadataConfiguration` request succeeds, but S3 Metadata was unable to create the table, this structure contains the error code and error message.

Note

If you created your S3 Metadata configuration before July 15, 2025, we recommend that you delete and re-create your configuration by using [CreateBucketMetadataConfiguration](#) so that you can expire journal table records and create a live inventory table.

Type: [ErrorDetails](#) data type

Required: No

TableArn

The Amazon Resource Name (ARN) for the journal table.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

JournalTableConfigurationUpdates

Service: Amazon S3

The specified updates to the S3 Metadata journal table configuration.

Contents

RecordExpiration

The journal table record expiration settings for the journal table.

Type: [RecordExpiration](#) data type

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

JSONInput

Service: Amazon S3

Specifies JSON as object's input serialization format.

Contents

Type

The type of JSON. Valid values: Document, Lines.

Type: String

Valid Values: DOCUMENT | LINES

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

JSONOutput

Service: Amazon S3

Specifies JSON as request's output serialization format.

Contents

RecordDelimiter

The value used to separate individual records in the output. If no value is specified, Amazon S3 uses a newline character ('\n').

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

LambdaFunctionConfiguration

Service: Amazon S3

A container for specifying the configuration for AWS Lambda notifications.

Contents

Events

The Amazon S3 bucket event for which to invoke the AWS Lambda function. For more information, see [Supported Event Types](#) in the *Amazon S3 User Guide*.

Type: Array of strings

Valid Values: s3:ReducedRedundancyLostObject | s3:ObjectCreated:* | s3:ObjectCreated:Put | s3:ObjectCreated:Post | s3:ObjectCreated:Copy | s3:ObjectCreated:CompleteMultipartUpload | s3:ObjectRemoved:* | s3:ObjectRemoved:Delete | s3:ObjectRemoved:DeleteMarkerCreated | s3:ObjectRestore:* | s3:ObjectRestore:Post | s3:ObjectRestore:Completed | s3:Replication:* | s3:Replication:OperationFailedReplication | s3:Replication:OperationNotTracked | s3:Replication:OperationMissedThreshold | s3:Replication:OperationReplicatedAfterThreshold | s3:ObjectRestore:Delete | s3:LifecycleTransition | s3: IntelligentTiering | s3:ObjectAcl:Put | s3:LifecycleExpiration:* | s3:LifecycleExpiration:Delete | s3:LifecycleExpiration:DeleteMarkerCreated | s3:ObjectTagging:* | s3:ObjectTagging:Put | s3:ObjectTagging:Delete | s3:ObjectAnnotation:* | s3:ObjectAnnotation:Put | s3:ObjectAnnotation:Delete | s3:ObjectRetention:Put

Required: Yes

LambdaFunctionArn

The Amazon Resource Name (ARN) of the AWS Lambda function that Amazon S3 invokes when the specified event type occurs.

Type: String

Required: Yes

Filter

Specifies object key name filtering rules. For information about key name filtering, see [Configuring event notifications using object key name filtering](#) in the *Amazon S3 User Guide*.

Type: [NotificationConfigurationFilter](#) data type

Required: No

Id

An optional unique identifier for configurations in a notification configuration. If you don't provide one, Amazon S3 will assign an ID.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

LifecycleConfiguration

Service: Amazon S3

Container for lifecycle rules. You can add as many as 1000 rules.

For more information see, [Managing your storage lifecycle](#) in the *Amazon S3 User Guide*.

Contents

Rules

Specifies lifecycle configuration rules for an Amazon S3 bucket.

Type: Array of [Rule](#) data types

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

LifecycleExpiration

Service: Amazon S3

Container for the expiration for the lifecycle of the object.

For more information see, [Managing your storage lifecycle](#) in the *Amazon S3 User Guide*.

Contents

Date

Indicates at what date the object is to be moved or deleted. The date value must conform to the ISO 8601 format. The time is always midnight UTC.

Note

This parameter applies to general purpose buckets only. It is not supported for directory bucket lifecycle configurations.

Type: Timestamp

Required: No

Days

Indicates the lifetime, in days, of the objects that are subject to the rule. The value must be a non-zero positive integer.

Type: Integer

Required: No

ExpiredObjectDeleteMarker

Indicates whether Amazon S3 will remove a delete marker with no noncurrent versions. If set to true, the delete marker will be expired; if set to false the policy takes no action. This cannot be specified with Days or Date in a Lifecycle Expiration Policy.

Note

This parameter applies to general purpose buckets only. It is not supported for directory bucket lifecycle configurations.

Type: Boolean

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

LifecycleRule

Service: Amazon S3

A lifecycle rule for individual objects in an Amazon S3 bucket.

For more information see, [Managing your storage lifecycle](#) in the *Amazon S3 User Guide*.

Contents

Status

If 'Enabled', the rule is currently being applied. If 'Disabled', the rule is not currently being applied.

Type: String

Valid Values: Enabled | Disabled

Required: Yes

AbortIncompleteMultipartUpload

Specifies the days since the initiation of an incomplete multipart upload that Amazon S3 will wait before permanently removing all parts of the upload. For more information, see [Aborting Incomplete Multipart Uploads Using a Bucket Lifecycle Configuration](#) in the *Amazon S3 User Guide*.

Type: [AbortIncompleteMultipartUpload](#) data type

Required: No

Expiration

Specifies the expiration for the lifecycle of the object in the form of date, days and, whether the object has a delete marker.

Type: [LifecycleExpiration](#) data type

Required: No

Filter

The `Filter` is used to identify objects that a Lifecycle Rule applies to. A `Filter` must have exactly one of `Prefix`, `Tag`, `ObjectSizeGreaterThan`, `ObjectSizeLessThan`, or `And` specified. `Filter` is required if the `LifecycleRule` does not contain a `Prefix` element.

For more information about Tag filters, see [Adding filters to Lifecycle rules](#) in the *Amazon S3 User Guide*.

Note

Tag filters are not supported for directory buckets.

Type: [LifecycleRuleFilter](#) data type

Required: No

ID

Unique identifier for the rule. The value cannot be longer than 255 characters.

Type: String

Required: No

NoncurrentVersionExpiration

Specifies when noncurrent object versions expire. Upon expiration, Amazon S3 permanently deletes the noncurrent object versions. You set this lifecycle configuration action on a bucket that has versioning enabled (or suspended) to request that Amazon S3 delete noncurrent object versions at a specific period in the object's lifetime.

Note

This parameter applies to general purpose buckets only. It is not supported for directory bucket lifecycle configurations.

Type: [NoncurrentVersionExpiration](#) data type

Required: No

NoncurrentVersionTransitions

Specifies the transition rule for the lifecycle rule that describes when noncurrent objects transition to a specific storage class. If your bucket is versioning-enabled (or versioning is suspended), you can set this action to request that Amazon S3 transition noncurrent object versions to a specific storage class at a set period in the object's lifetime.

Note

This parameter applies to general purpose buckets only. It is not supported for directory bucket lifecycle configurations.

Type: Array of [NoncurrentVersionTransition](#) data types

Required: No

Prefix

This member has been deprecated.

The general purpose bucket prefix that identifies one or more objects to which the rule applies. We recommend using `Filter` instead of `Prefix` for new PUTs. Previous configurations where a prefix is defined will continue to operate as before.

Important

Replacement must be made for object keys containing special characters (such as carriage returns) when using XML requests. For more information, see [XML related object key constraints](#).

Type: String

Required: No

Transitions

Specifies when an Amazon S3 object transitions to a specified storage class.

Note

This parameter applies to general purpose buckets only. It is not supported for directory bucket lifecycle configurations.

Type: Array of [Transition](#) data types

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

LifecycleRuleAndOperator

Service: Amazon S3

This is used in a Lifecycle Rule Filter to apply a logical AND to two or more predicates. The Lifecycle Rule will apply to any object matching all of the predicates configured inside the And operator.

Contents

ObjectSizeGreaterThan

Minimum object size to which the rule applies.

Type: Long

Required: No

ObjectSizeLessThan

Maximum object size to which the rule applies.

Type: Long

Required: No

Prefix

Prefix identifying one or more objects to which the rule applies.

Type: String

Required: No

Tags

All of these tags must exist in the object's tag set in order for the rule to apply.

Type: Array of [Tag](#) data types

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

LifecycleRuleFilter

Service: Amazon S3

The `Filter` is used to identify objects that a Lifecycle Rule applies to. A `Filter` can have exactly one of `Prefix`, `Tag`, `ObjectSizeGreaterThan`, `ObjectSizeLessThan`, or `And` specified. If the `Filter` element is left empty, the Lifecycle Rule applies to all objects in the bucket.

Contents

And

This is used in a Lifecycle Rule Filter to apply a logical AND to two or more predicates. The Lifecycle Rule will apply to any object matching all of the predicates configured inside the `And` operator.

Type: [LifecycleRuleAndOperator](#) data type

Required: No

ObjectSizeGreaterThan

Minimum object size to which the rule applies.

Type: Long

Required: No

ObjectSizeLessThan

Maximum object size to which the rule applies.

Type: Long

Required: No

Prefix

Prefix identifying one or more objects to which the rule applies.

Important

Replacement must be made for object keys containing special characters (such as carriage returns) when using XML requests. For more information, see [XML related object key constraints](#).

Type: String

Required: No

Tag

This tag must exist in the object's tag set in order for the rule to apply.

Note

This parameter applies to general purpose buckets only. It is not supported for directory bucket lifecycle configurations.

Type: [Tag](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

LocationInfo

Service: Amazon S3

Specifies the location where the bucket will be created.

For directory buckets, the location type is Availability Zone or Local Zone. For more information about directory buckets, see [Working with directory buckets](#) in the *Amazon S3 User Guide*.

Note

This functionality is only supported by directory buckets.

Contents

Name

The name of the location where the bucket will be created.

For directory buckets, the name of the location is the Zone ID of the Availability Zone (AZ) or Local Zone (LZ) where the bucket will be created. An example AZ ID value is `usw2-az1`.

Type: String

Required: No

Type

The type of location where the bucket will be created.

Type: String

Valid Values: `AvailabilityZone` | `LocalZone`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

LoggingEnabled

Service: Amazon S3

Describes where logs are stored and the prefix that Amazon S3 assigns to all log object keys for a bucket. For more information, see [PUT Bucket logging](#) in the *Amazon S3 API Reference*.

Contents

TargetBucket

Specifies the bucket where you want Amazon S3 to store server access logs. You can have your logs delivered to any bucket that you own, including the same bucket that is being logged. You can also configure multiple buckets to deliver their logs to the same target bucket. In this case, you should choose a different `TargetPrefix` for each source bucket so that the delivered log files can be distinguished by key.

Type: String

Required: Yes

TargetPrefix

A prefix for all log object keys. If you store log files from multiple Amazon S3 buckets in a single bucket, you can use a prefix to distinguish which log files came from which bucket.

Type: String

Required: Yes

TargetGrants

Container for granting information.

Buckets that use the bucket owner enforced setting for Object Ownership don't support target grants. For more information, see [Permissions for server access log delivery](#) in the *Amazon S3 User Guide*.

Type: Array of [TargetGrant](#) data types

Required: No

TargetObjectKeyFormat

Amazon S3 key format for log objects.

Type: [TargetObjectKeyFormat](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

MetadataConfiguration

Service: Amazon S3

The S3 Metadata configuration for a general purpose bucket.

Contents

JournalTableConfiguration

The journal table configuration for a metadata configuration.

Type: [JournalTableConfiguration](#) data type

Required: Yes

AnnotationTableConfiguration

Optional annotation table configuration to include with the metadata configuration.

Type: [AnnotationTableConfiguration](#) data type

Required: No

InventoryTableConfiguration

The inventory table configuration for a metadata configuration.

Type: [InventoryTableConfiguration](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

MetadataConfigurationResult

Service: Amazon S3

The S3 Metadata configuration for a general purpose bucket.

Contents

DestinationResult

The destination settings for a metadata configuration.

Type: [DestinationResult](#) data type

Required: Yes

AnnotationTableConfigurationResult

The annotation table configuration result, if an annotation table is configured.

Type: [AnnotationTableConfigurationResult](#) data type

Required: No

InventoryTableConfigurationResult

The inventory table configuration for a metadata configuration.

Type: [InventoryTableConfigurationResult](#) data type

Required: No

JournalTableConfigurationResult

The journal table configuration for a metadata configuration.

Type: [JournalTableConfigurationResult](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

MetadataEntry

Service: Amazon S3

A metadata key-value pair to store with an object.

Contents

Name

Name of the object.

Type: String

Required: No

Value

Value of the object.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

MetadataTableConfiguration

Service: Amazon S3

The V1 S3 Metadata configuration for a general purpose bucket.

Note

If you created your S3 Metadata configuration before July 15, 2025, we recommend that you delete and re-create your configuration by using [CreateBucketMetadataConfiguration](#) so that you can expire journal table records and create a live inventory table.

Contents

S3TablesDestination

The destination information for the metadata table configuration. The destination table bucket must be in the same Region and AWS account as the general purpose bucket. The specified metadata table name must be unique within the `aws_s3_metadata` namespace in the destination table bucket.

Type: [S3TablesDestination](#) data type

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

MetadataTableConfigurationResult

Service: Amazon S3

The V1 S3 Metadata configuration for a general purpose bucket. The destination table bucket must be in the same Region and AWS account as the general purpose bucket. The specified metadata table name must be unique within the `aws_s3_metadata` namespace in the destination table bucket.

Note

If you created your S3 Metadata configuration before July 15, 2025, we recommend that you delete and re-create your configuration by using [CreateBucketMetadataConfiguration](#) so that you can expire journal table records and create a live inventory table.

Contents

S3TablesDestinationResult

The destination information for the metadata table configuration. The destination table bucket must be in the same Region and AWS account as the general purpose bucket. The specified metadata table name must be unique within the `aws_s3_metadata` namespace in the destination table bucket.

Type: [S3TablesDestinationResult](#) data type

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

MetadataTableEncryptionConfiguration

Service: Amazon S3

The encryption settings for an S3 Metadata journal table or inventory table configuration.

Contents

SseAlgorithm

The encryption type specified for a metadata table. To specify server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS), use the `aws:kms` value. To specify server-side encryption with Amazon S3 managed keys (SSE-S3), use the `AES256` value.

Type: String

Valid Values: `aws:kms` | `AES256`

Required: Yes

KmsKeyArn

If server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS) is specified, you must also specify the KMS key Amazon Resource Name (ARN). You must specify a customer-managed KMS key that's located in the same Region as the general purpose bucket that corresponds to the metadata table configuration.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Metrics

Service: Amazon S3

A container specifying replication metrics-related settings enabling replication metrics and events.

Contents

Status

Specifies whether the replication metrics are enabled.

Type: String

Valid Values: Enabled | Disabled

Required: Yes

EventThreshold

A container specifying the time threshold for emitting the `s3:Replication:OperationMissedThreshold` event.

Type: [ReplicationTimeValue](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

MetricsAndOperator

Service: Amazon S3

A conjunction (logical AND) of predicates, which is used in evaluating a metrics filter. The operator must have at least two predicates, and an object must match all of the predicates in order for the filter to apply.

Contents

AccessPointArn

The access point ARN used when evaluating an AND predicate.

Type: String

Required: No

Prefix

The prefix used when evaluating an AND predicate.

Type: String

Required: No

Tags

The list of tags used when evaluating an AND predicate.

Note

Tag filters are not supported for directory buckets.

Type: Array of [Tag](#) data types

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

MetricsConfiguration

Service: Amazon S3

Specifies a metrics configuration for the CloudWatch request metrics (specified by the metrics configuration ID) from an Amazon S3 bucket. If you're updating an existing metrics configuration, note that this is a full replacement of the existing metrics configuration. If you don't include the elements you want to keep, they are erased. For more information, see [PutBucketMetricsConfiguration](#).

Contents

Id

The ID used to identify the metrics configuration. The ID has a 64 character limit and can only contain letters, numbers, periods, dashes, and underscores.

Type: String

Required: Yes

Filter

Specifies a metrics configuration filter. The metrics configuration will only include objects that meet the filter's criteria. A filter must be a prefix, an object tag, an access point ARN, or a conjunction (MetricsAndOperator).

Note

Metrics configurations for directory buckets do not support tag filters.

Type: [MetricsFilter](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

MetricsFilter

Service: Amazon S3

Specifies a metrics configuration filter. The metrics configuration only includes objects that meet the filter's criteria. A filter must be a prefix, an object tag, an access point ARN, or a conjunction (MetricsAndOperator). For more information, see [PutBucketMetricsConfiguration](#).

Contents

AccessPointArn

The access point ARN used when evaluating a metrics filter.

Type: String

Required: No

And

A conjunction (logical AND) of predicates, which is used in evaluating a metrics filter. The operator must have at least two predicates, and an object must match all of the predicates in order for the filter to apply.

Type: [MetricsAndOperator](#) data type

Required: No

Prefix

The prefix used when evaluating a metrics filter.

Type: String

Required: No

Tag

The tag used when evaluating a metrics filter.

Note

Tag filters are not supported for directory buckets.

Type: [Tag](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

MultipartUpload

Service: Amazon S3

Container for the MultipartUpload for the Amazon S3 object.

Contents

ChecksumAlgorithm

The algorithm that was used to create a checksum of the object.

Type: String

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 |
XXHASH64 | XXHASH3 | XXHASH128

Required: No

ChecksumType

The checksum type that is used to calculate the object's checksum value. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Valid Values: COMPOSITE | FULL_OBJECT

Required: No

Initiated

Date and time at which the multipart upload was initiated.

Type: Timestamp

Required: No

Initiator

Identifies who initiated the multipart upload.

Type: [Initiator](#) data type

Required: No

Key

Key of the object for which the multipart upload was initiated.

Type: String

Length Constraints: Minimum length of 1.

Required: No

Owner

Specifies the owner of the object that is part of the multipart upload.

Note

Directory buckets - The bucket owner is returned as the object owner for all the objects.

Type: [Owner](#) data type

Required: No

StorageClass

The class of storage used to store the object.

Note

Directory buckets - Directory buckets only support EXPRESS_ONEZONE (the S3 Express One Zone storage class) in Availability Zones and ONEZONE_IA (the S3 One Zone-Infrequent Access storage class) in Dedicated Local Zones.

Type: String

Valid Values: STANDARD | REDUCED_REDUNDANCY | STANDARD_IA | ONEZONE_IA | INTELLIGENT_TIERING | GLACIER | DEEP_ARCHIVE | OUTPOSTS | GLACIER_IR | SNOW | EXPRESS_ONEZONE | FSX_OPENZFS | FSX_ONTAP | AWS_BACKUP_WARM | AWS_BACKUP_LOW_COST_WARM

Required: No

UploadId

Upload ID that identifies the multipart upload.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

NoncurrentVersionExpiration

Service: Amazon S3

Specifies when noncurrent object versions expire. Upon expiration, Amazon S3 permanently deletes the noncurrent object versions. You set this lifecycle configuration action on a bucket that has versioning enabled (or suspended) to request that Amazon S3 delete noncurrent object versions at a specific period in the object's lifetime.

Note

This parameter applies to general purpose buckets only. It is not supported for directory bucket lifecycle configurations.

Contents

NewerNoncurrentVersions

Specifies how many noncurrent versions Amazon S3 will retain. You can specify up to 100 noncurrent versions to retain. Amazon S3 will permanently delete any additional noncurrent versions beyond the specified number to retain. For more information about noncurrent versions, see [Lifecycle configuration elements](#) in the *Amazon S3 User Guide*.

Note

This parameter applies to general purpose buckets only. It is not supported for directory bucket lifecycle configurations.

Type: Integer

Required: No

NoncurrentDays

Specifies the number of days an object is noncurrent before Amazon S3 can perform the associated action. The value must be a non-zero positive integer. For information about the noncurrent days calculations, see [How Amazon S3 Calculates When an Object Became Noncurrent](#) in the *Amazon S3 User Guide*.

Note

This parameter applies to general purpose buckets only. It is not supported for directory bucket lifecycle configurations.

Type: Integer

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

NoncurrentVersionTransition

Service: Amazon S3

Container for the transition rule that describes when noncurrent objects transition to the STANDARD_IA, ONEZONE_IA, INTELLIGENT_TIERING, GLACIER_IR, GLACIER, or DEEP_ARCHIVE storage class. If your bucket is versioning-enabled (or versioning is suspended), you can set this action to request that Amazon S3 transition noncurrent object versions to the STANDARD_IA, ONEZONE_IA, INTELLIGENT_TIERING, GLACIER_IR, GLACIER, or DEEP_ARCHIVE storage class at a specific period in the object's lifetime.

Contents

NewerNoncurrentVersions

Specifies how many noncurrent versions Amazon S3 will retain in the same storage class before transitioning objects. You can specify up to 100 noncurrent versions to retain. Amazon S3 will transition any additional noncurrent versions beyond the specified number to retain. For more information about noncurrent versions, see [Lifecycle configuration elements](#) in the *Amazon S3 User Guide*.

Type: Integer

Required: No

NoncurrentDays

Specifies the number of days an object is noncurrent before Amazon S3 can perform the associated action. For information about the noncurrent days calculations, see [How Amazon S3 Calculates How Long an Object Has Been Noncurrent](#) in the *Amazon S3 User Guide*.

Type: Integer

Required: No

StorageClass

The class of storage used to store the object.

Type: String

Valid Values: GLACIER | STANDARD_IA | ONEZONE_IA | INTELLIGENT_TIERING | DEEP_ARCHIVE | GLACIER_IR

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

NotificationConfiguration

Service: Amazon S3

A container for specifying the notification configuration of the bucket. If this element is empty, notifications are turned off for the bucket.

Contents

EventBridgeConfiguration

Enables delivery of events to Amazon EventBridge.

Type: [EventBridgeConfiguration](#) data type

Required: No

LambdaFunctionConfigurations

Describes the AWS Lambda functions to invoke and the events for which to invoke them.

Type: Array of [LambdaFunctionConfiguration](#) data types

Required: No

QueueConfigurations

The Amazon Simple Queue Service queues to publish messages to and the events for which to publish messages.

Type: Array of [QueueConfiguration](#) data types

Required: No

TopicConfigurations

The topic to which notifications are sent and the events for which notifications are generated.

Type: Array of [TopicConfiguration](#) data types

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

NotificationConfigurationDeprecated

Service: Amazon S3

Contents

CloudFunctionConfiguration

Container for specifying the AWS Lambda notification configuration.

Type: [CloudFunctionConfiguration](#) data type

Required: No

QueueConfiguration

This data type is deprecated. This data type specifies the configuration for publishing messages to an Amazon Simple Queue Service (Amazon SQS) queue when Amazon S3 detects specified events.

Type: [QueueConfigurationDeprecated](#) data type

Required: No

TopicConfiguration

This data type is deprecated. A container for specifying the configuration for publication of messages to an Amazon Simple Notification Service (Amazon SNS) topic when Amazon S3 detects specified events.

Type: [TopicConfigurationDeprecated](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

NotificationConfigurationFilter

Service: Amazon S3

Specifies object key name filtering rules. For information about key name filtering, see [Configuring event notifications using object key name filtering](#) in the *Amazon S3 User Guide*.

Contents

Key

A container for object key name prefix and suffix filtering rules.

Type: [S3KeyFilter](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Object

Service: Amazon S3

An object consists of data and its descriptive metadata.

Contents

ChecksumAlgorithm

The algorithm that was used to create a checksum of the object.

Type: Array of strings

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 |
XXHASH64 | XXHASH3 | XXHASH128

Required: No

ChecksumType

The checksum type that is used to calculate the object's checksum value. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Valid Values: COMPOSITE | FULL_OBJECT

Required: No

ETag

The entity tag is a hash of the object. The ETag reflects changes only to the contents of an object, not its metadata. The ETag may or may not be an MD5 digest of the object data. Whether or not it is depends on how the object was created and how it is encrypted as described below:

- Objects created by the PUT Object, POST Object, or Copy operation, or through the AWS Management Console, and are encrypted by SSE-S3 or plaintext, have ETags that are an MD5 digest of their object data.
- Objects created by the PUT Object, POST Object, or Copy operation, or through the AWS Management Console, and are encrypted by SSE-C or SSE-KMS, have ETags that are not an MD5 digest of their object data.

- If an object is created by either the Multipart Upload or Part Copy operation, the ETag is not an MD5 digest, regardless of the method of encryption. If an object is larger than 16 MB, the AWS Management Console will upload or copy that object as a Multipart Upload, and therefore the ETag will not be an MD5 digest.

Note

Directory buckets - MD5 is not supported by directory buckets.

Type: String

Required: No

Key

The name that you assign to an object. You use the object key to retrieve the object.

Type: String

Length Constraints: Minimum length of 1.

Required: No

LastModified

Creation date of the object.

Type: Timestamp

Required: No

Owner

The owner of the object

Note

Directory buckets - The bucket owner is returned as the object owner.

Type: [Owner](#) data type

Required: No

RestoreStatus

Specifies the restoration status of an object. Objects in certain storage classes must be restored before they can be retrieved. For more information about these storage classes and how to work with archived objects, see [Working with archived objects](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets. Directory buckets only support EXPRESS_ONEZONE (the S3 Express One Zone storage class) in Availability Zones and ONEZONE_IA (the S3 One Zone-Infrequent Access storage class) in Dedicated Local Zones.

Type: [RestoreStatus](#) data type

Required: No

Size

Size in bytes of the object

Type: Long

Required: No

StorageClass

The class of storage used to store the object.

Note

Directory buckets - Directory buckets only support EXPRESS_ONEZONE (the S3 Express One Zone storage class) in Availability Zones and ONEZONE_IA (the S3 One Zone-Infrequent Access storage class) in Dedicated Local Zones.

Type: String

Valid Values: STANDARD | REDUCED_REDUNDANCY | GLACIER | STANDARD_IA | ONEZONE_IA | INTELLIGENT_TIERING | DEEP_ARCHIVE | OUTPOSTS | GLACIER_IR

| SNOW | EXPRESS_ONEZONE | FSX_OPENZFS | FSX_ONTAP | AWS_BACKUP_WARM |
AWS_BACKUP_LOW_COST_WARM

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ObjectEncryption

Service: Amazon S3

The updated server-side encryption type for this object. The `UpdateObjectEncryption` operation supports the SSE-S3 and SSE-KMS encryption types.

Valid Values: SSE3 | SSEKMS

Contents

Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

SSEKMS

Specifies to update the object encryption type to server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS).

Type: [SSEKMSEncryption](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ObjectIdentifier

Service: Amazon S3

Object Identifier is unique value to identify objects.

Contents

Key

Key name of the object.

Important

Replacement must be made for object keys containing special characters (such as carriage returns) when using XML requests. For more information, see [XML related object key constraints](#).

Type: String

Length Constraints: Minimum length of 1.

Required: Yes

ETag

An entity tag (ETag) is an identifier assigned by a web server to a specific version of a resource found at a URL. This header field makes the request method conditional on ETags.

Note

Entity tags (ETags) for S3 Express One Zone are random alphanumeric strings unique to the object.

Type: String

Required: No

LastModifiedTime

If present, the objects are deleted only if its modification times matches the provided Timestamp.

Note

This functionality is only supported for directory buckets.

Type: Timestamp

Required: No

Size

If present, the objects are deleted only if its size matches the provided size in bytes.

Note

This functionality is only supported for directory buckets.

Type: Long

Required: No

VersionId

Version ID for the specific version of the object to delete.

Note

This functionality is not supported for directory buckets.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ObjectLockConfiguration

Service: Amazon S3

The container element for Object Lock configuration parameters.

Contents

ObjectLockEnabled

Indicates whether this bucket has an Object Lock configuration enabled. Enable `ObjectLockEnabled` when you apply `ObjectLockConfiguration` to a bucket.

Type: String

Valid Values: Enabled

Required: No

Rule

Specifies the Object Lock rule for the specified object. Enable the this rule when you apply `ObjectLockConfiguration` to a bucket. Bucket settings require both a mode and a period. The period can be either `Days` or `Years` but you must select one. You cannot specify `Days` and `Years` at the same time.

Type: [ObjectLockRule](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ObjectLockLegalHold

Service: Amazon S3

A legal hold configuration for an object.

Contents

Status

Indicates whether the specified object has a legal hold in place.

Type: String

Valid Values: ON | OFF

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ObjectLockRetention

Service: Amazon S3

A Retention configuration for an object.

Contents

EventHold

The event hold status for the object. Set to ON to enable an event hold or OFF to disable it.

Type: String

Valid Values: ON | OFF

Required: No

EventHoldDuration

The event hold duration for the object. Specifies how long the object remains protected after the event hold is released.

Type: [EventHoldDuration](#) data type

Required: No

Mode

Indicates the Retention mode for the specified object.

Type: String

Valid Values: GOVERNANCE | COMPLIANCE

Required: No

RetainUntilDate

The date on which this Object Lock Retention will expire.

Type: Timestamp

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ObjectLockRule

Service: Amazon S3

The container element for an Object Lock rule.

Contents

DefaultRetention

The default Object Lock retention settings for new objects in this bucket. You can specify:

- A default retention period, by using `Days` or `Years`.
- A default event hold duration, by using `DefaultEventHold`. This setting also uses days or years.

You can set one or both. You cannot use days and years in the same setting.

Type: [DefaultRetention](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ObjectPart

Service: Amazon S3

A container for elements related to an individual part.

Contents

ChecksumCRC32

The Base64 encoded, 32-bit CRC32 checksum of the part. This checksum is present if the multipart upload request was created with the CRC32 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumCRC32C

The Base64 encoded, 32-bit CRC32C checksum of the part. This checksum is present if the multipart upload request was created with the CRC32C checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumCRC64NVME

The Base64 encoded, 64-bit CRC64NVME checksum of the part. This checksum is present if the multipart upload request was created with the CRC64NVME checksum algorithm, or if the object was uploaded without a checksum (and Amazon S3 added the default checksum, CRC64NVME, to the uploaded object). For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumMD5

The Base64 encoded, 128-bit MD5 digest of the part. This checksum is present if the multipart upload request was created with the MD5 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumSHA1

The Base64 encoded, 160-bit SHA1 checksum of the part. This checksum is present if the multipart upload request was created with the SHA1 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumSHA256

The Base64 encoded, 256-bit SHA256 checksum of the part. This checksum is present if the multipart upload request was created with the SHA256 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumSHA512

The Base64 encoded, 512-bit SHA512 digest of the part. This checksum is present if the multipart upload request was created with the SHA512 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumXXHASH128

The Base64 encoded, 128-bit XXHASH128 checksum of the part. This checksum is present if the multipart upload request was created with the XXHASH128 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumXXHASH3

The Base64 encoded, 64-bit XXHASH3 checksum of the part. This checksum is present if the multipart upload request was created with the XXHASH3 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumXXHASH64

The Base64 encoded, 64-bit XXHASH64 checksum of the part. This checksum is present if the multipart upload request was created with the XXHASH64 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

PartNumber

The part number identifying the part. This value is a positive integer between 1 and 10,000.

Type: Integer

Required: No

Size

The size of the uploaded part in bytes.

Type: Long

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)

- [AWS SDK for Ruby V3](#)

ObjectVersion

Service: Amazon S3

The version of an object.

Contents

ChecksumAlgorithm

The algorithm that was used to create a checksum of the object.

Type: Array of strings

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 |
XXHASH64 | XXHASH3 | XXHASH128

Required: No

ChecksumType

The checksum type that is used to calculate the object's checksum value. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Valid Values: COMPOSITE | FULL_OBJECT

Required: No

ETag

The entity tag is an MD5 hash of that version of the object.

Type: String

Required: No

IsLatest

Specifies whether the object is (true) or is not (false) the latest version of an object.

Type: Boolean

Required: No

Key

The object key.

Type: String

Length Constraints: Minimum length of 1.

Required: No

LastModified

Date and time when the object was last modified.

Type: Timestamp

Required: No

Owner

Specifies the owner of the object.

Type: [Owner](#) data type

Required: No

RestoreStatus

Specifies the restoration status of an object. Objects in certain storage classes must be restored before they can be retrieved. For more information about these storage classes and how to work with archived objects, see [Working with archived objects](#) in the *Amazon S3 User Guide*.

Type: [RestoreStatus](#) data type

Required: No

Size

Size in bytes of the object.

Type: Long

Required: No

StorageClass

The class of storage used to store the object.

Type: String

Valid Values: STANDARD

Required: No

VersionId

Version ID of an object.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

OutputLocation

Service: Amazon S3

Describes the location where the restore job's output is stored.

Contents

S3

Describes an S3 location that will receive the results of the restore request.

Type: [S3Location](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

OutputSerialization

Service: Amazon S3

Describes how results of the Select job are serialized.

Contents

CSV

Describes the serialization of CSV-encoded Select results.

Type: [CSVOutput](#) data type

Required: No

JSON

Specifies JSON as request's output serialization format.

Type: [JSONOutput](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Owner

Service: Amazon S3

Container for the owner's display name and ID.

Contents

DisplayName

Type: String

Required: No

ID

Container for the ID of the owner.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

OwnershipControls

Service: Amazon S3

The container element for a bucket's ownership controls.

Contents

Rules

The container element for an ownership control rule.

Type: Array of [OwnershipControlsRule](#) data types

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

OwnershipControlsRule

Service: Amazon S3

The container element for an ownership control rule.

Contents

ObjectOwnership

The container element for object ownership for a bucket's ownership controls.

`BucketOwnerPreferred` - Objects uploaded to the bucket change ownership to the bucket owner if the objects are uploaded with the `bucket-owner-full-control` canned ACL.

`ObjectWriter` - The uploading account will own the object if the object is uploaded with the `bucket-owner-full-control` canned ACL.

`BucketOwnerEnforced` - Access control lists (ACLs) are disabled and no longer affect permissions. The bucket owner automatically owns and has full control over every object in the bucket. The bucket only accepts PUT requests that don't specify an ACL or specify bucket owner full control ACLs (such as the predefined `bucket-owner-full-control` canned ACL or a custom ACL in XML format that grants the same permissions).

By default, `ObjectOwnership` is set to `BucketOwnerEnforced` and ACLs are disabled. We recommend keeping ACLs disabled, except in uncommon use cases where you must control access for each object individually. For more information about S3 Object Ownership, see [Controlling ownership of objects and disabling ACLs for your bucket](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets. Directory buckets use the bucket owner enforced setting for S3 Object Ownership.

Type: String

Valid Values: `BucketOwnerPreferred` | `ObjectWriter` | `BucketOwnerEnforced`

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ParquetInput

Service: Amazon S3

Container for Parquet.

Contents

The members of this exception structure are context-dependent.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Part

Service: Amazon S3

Container for elements related to a part.

Contents

ChecksumCRC32

The Base64 encoded, 32-bit CRC32 checksum of the part. This checksum is present if the object was uploaded with the CRC32 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumCRC32C

The Base64 encoded, 32-bit CRC32C checksum of the part. This checksum is present if the object was uploaded with the CRC32C checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumCRC64NVME

The Base64 encoded, 64-bit CRC64NVME checksum of the part. This checksum is present if the multipart upload request was created with the CRC64NVME checksum algorithm, or if the object was uploaded without a checksum (and Amazon S3 added the default checksum, CRC64NVME, to the uploaded object). For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumMD5

The Base64 encoded, 128-bit MD5 digest of the part. This checksum is present if the multipart upload request was created with the MD5 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumSHA1

The Base64 encoded, 160-bit SHA1 checksum of the part. This checksum is present if the object was uploaded with the SHA1 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumSHA256

The Base64 encoded, 256-bit SHA256 checksum of the part. This checksum is present if the object was uploaded with the SHA256 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumSHA512

The Base64 encoded, 512-bit SHA512 digest of the part. This checksum is present if the multipart upload request was created with the SHA512 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumXXHASH128

The Base64 encoded, 128-bit XXHASH128 checksum of the part. This checksum is present if the multipart upload request was created with the XXHASH128 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumXXHASH3

The Base64 encoded, 64-bit XXHASH3 checksum of the part. This checksum is present if the multipart upload request was created with the XXHASH3 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ChecksumXXHASH64

The Base64 encoded, 64-bit XXHASH64 checksum of the part. This checksum is present if the multipart upload request was created with the XXHASH64 checksum algorithm. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

ETag

Entity tag returned when the part was uploaded.

Type: String

Required: No

LastModified

Date and time at which the part was uploaded.

Type: Timestamp

Required: No

PartNumber

Part number identifying the part. This is a positive integer between 1 and 10,000.

Type: Integer

Required: No

Size

Size in bytes of the uploaded part data.

Type: Long

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

PartitionedPrefix

Service: Amazon S3

Amazon S3 keys for log objects are partitioned in the following format:

```
[DestinationPrefix][SourceAccountId]/[SourceRegion]/[SourceBucket]/[YYYY]/[MM]/[DD]/[YYYY]-[MM]-[DD]-[hh]-[mm]-[ss]-[UniqueString]
```

PartitionedPrefix defaults to EventTime delivery when server access logs are delivered.

Contents

PartitionDateSource

Specifies the partition date source for the partitioned prefix. PartitionDateSource can be EventTime or DeliveryTime.

For DeliveryTime, the time in the log file names corresponds to the delivery time for the log files.

For EventTime, The logs delivered are for a specific day only. The year, month, and day correspond to the day on which the event occurred, and the hour, minutes and seconds are set to 00 in the key.

Type: String

Valid Values: EventTime | DeliveryTime

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

PolicyStatus

Service: Amazon S3

The container element for a bucket's policy status.

Contents

IsPublic

The policy status for this bucket. TRUE indicates that this bucket is public. FALSE indicates that the bucket is not public.

Type: Boolean

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Progress

Service: Amazon S3

This data type contains information about progress of an operation.

Contents

BytesProcessed

The current number of uncompressed object bytes processed.

Type: Long

Required: No

BytesReturned

The current number of bytes of records payload data returned.

Type: Long

Required: No

BytesScanned

The current number of object bytes scanned.

Type: Long

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ProgressEvent

Service: Amazon S3

This data type contains information about the progress event of an operation.

Contents

Details

The Progress event details.

Type: [Progress](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

PublicAccessBlockConfiguration

Service: Amazon S3

The PublicAccessBlock configuration that you want to apply to this Amazon S3 bucket. You can enable the configuration options in any combination. Bucket-level settings work alongside account-level settings (which may inherit from organization-level policies). For more information about when Amazon S3 considers a bucket or object public, see [The Meaning of "Public"](#) in the *Amazon S3 User Guide*.

Contents

BlockPublicAcls

Specifies whether Amazon S3 should block public access control lists (ACLs) for this bucket and objects in this bucket. Setting this element to TRUE causes the following behavior:

- PUT Bucket ACL and PUT Object ACL calls fail if the specified ACL is public.
- PUT Object calls fail if the request includes a public ACL.
- PUT Bucket calls fail if the request includes a public ACL.

Enabling this setting doesn't affect existing policies or ACLs.

Type: Boolean

Required: No

BlockPublicPolicy

Specifies whether Amazon S3 should block public bucket policies for this bucket. Setting this element to TRUE causes Amazon S3 to reject calls to PUT Bucket policy if the specified bucket policy allows public access.

Enabling this setting doesn't affect existing bucket policies.

Type: Boolean

Required: No

IgnorePublicAcls

Specifies whether Amazon S3 should ignore public ACLs for this bucket and objects in this bucket. Setting this element to TRUE causes Amazon S3 to ignore all public ACLs on this bucket and objects in this bucket.

Enabling this setting doesn't affect the persistence of any existing ACLs and doesn't prevent new public ACLs from being set.

Type: Boolean

Required: No

RestrictPublicBuckets

Specifies whether Amazon S3 should restrict public bucket policies for this bucket. Setting this element to TRUE restricts access to this bucket to only AWS service principals and authorized users within this account if the bucket has a public policy.

Enabling this setting doesn't affect previously stored bucket policies, except that public and cross-account access within any public bucket policy, including non-public delegation to specific accounts, is blocked.

Type: Boolean

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

QueueConfiguration

Service: Amazon S3

Specifies the configuration for publishing messages to an Amazon Simple Queue Service (Amazon SQS) queue when Amazon S3 detects specified events.

Contents

Events

A collection of bucket events for which to send notifications

Type: Array of strings

Valid Values: `s3:ReducedRedundancyLostObject` | `s3:ObjectCreated:*` | `s3:ObjectCreated:Put` | `s3:ObjectCreated:Post` | `s3:ObjectCreated:Copy` | `s3:ObjectCreated:CompleteMultipartUpload` | `s3:ObjectRemoved:*` | `s3:ObjectRemoved:Delete` | `s3:ObjectRemoved:DeleteMarkerCreated` | `s3:ObjectRestore:*` | `s3:ObjectRestore:Post` | `s3:ObjectRestore:Completed` | `s3:Replication:*` | `s3:Replication:OperationFailedReplication` | `s3:Replication:OperationNotTracked` | `s3:Replication:OperationMissedThreshold` | `s3:Replication:OperationReplicatedAfterThreshold` | `s3:ObjectRestore:Delete` | `s3:LifecycleTransition` | `s3:IntelligentTiering` | `s3:ObjectAcl:Put` | `s3:LifecycleExpiration:*` | `s3:LifecycleExpiration:Delete` | `s3:LifecycleExpiration:DeleteMarkerCreated` | `s3:ObjectTagging:*` | `s3:ObjectTagging:Put` | `s3:ObjectTagging:Delete` | `s3:ObjectAnnotation:*` | `s3:ObjectAnnotation:Put` | `s3:ObjectAnnotation:Delete` | `s3:ObjectRetention:Put`

Required: Yes

QueueArn

The Amazon Resource Name (ARN) of the Amazon SQS queue to which Amazon S3 publishes a message when it detects events of the specified type.

Type: String

Required: Yes

Filter

Specifies object key name filtering rules. For information about key name filtering, see [Configuring event notifications using object key name filtering](#) in the *Amazon S3 User Guide*.

Type: [NotificationConfigurationFilter](#) data type

Required: No

Id

An optional unique identifier for configurations in a notification configuration. If you don't provide one, Amazon S3 will assign an ID.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

QueueConfigurationDeprecated

Service: Amazon S3

This data type is deprecated. Use [QueueConfiguration](#) for the same purposes. This data type specifies the configuration for publishing messages to an Amazon Simple Queue Service (Amazon SQS) queue when Amazon S3 detects specified events.

Contents

Event

This member has been deprecated.

The bucket event for which to send notifications.

Type: String

Valid Values: s3:ReducedRedundancyLostObject | s3:ObjectCreated:* | s3:ObjectCreated:Put | s3:ObjectCreated:Post | s3:ObjectCreated:Copy | s3:ObjectCreated:CompleteMultipartUpload | s3:ObjectRemoved:* | s3:ObjectRemoved:Delete | s3:ObjectRemoved:DeleteMarkerCreated | s3:ObjectRestore:* | s3:ObjectRestore:Post | s3:ObjectRestore:Completed | s3:Replication:* | s3:Replication:OperationFailedReplication | s3:Replication:OperationNotTracked | s3:Replication:OperationMissedThreshold | s3:Replication:OperationReplicatedAfterThreshold | s3:ObjectRestore:Delete | s3:LifecycleTransition | s3: IntelligentTiering | s3:ObjectAcl:Put | s3:LifecycleExpiration:* | s3:LifecycleExpiration:Delete | s3:LifecycleExpiration:DeleteMarkerCreated | s3:ObjectTagging:* | s3:ObjectTagging:Put | s3:ObjectTagging:Delete | s3:ObjectAnnotation:* | s3:ObjectAnnotation:Put | s3:ObjectAnnotation:Delete | s3:ObjectRetention:Put

Required: No

Events

A collection of bucket events for which to send notifications.

Type: Array of strings

Valid Values: s3:ReducedRedundancyLostObject | s3:ObjectCreated:* | s3:ObjectCreated:Put | s3:ObjectCreated:Post | s3:ObjectCreated:Copy | s3:ObjectCreated:CompleteMultipartUpload | s3:ObjectRemoved:* | s3:ObjectRemoved:Delete | s3:ObjectRemoved:DeleteMarkerCreated | s3:ObjectRestore:* | s3:ObjectRestore:Post | s3:ObjectRestore:Completed | s3:Replication:* | s3:Replication:OperationFailedReplication | s3:Replication:OperationNotTracked | s3:Replication:OperationMissedThreshold | s3:Replication:OperationReplicatedAfterThreshold | s3:ObjectRestore:Delete | s3:LifecycleTransition | s3:IntelligentTiering | s3:ObjectAcl:Put | s3:LifecycleExpiration:* | s3:LifecycleExpiration:Delete | s3:LifecycleExpiration:DeleteMarkerCreated | s3:ObjectTagging:* | s3:ObjectTagging:Put | s3:ObjectTagging:Delete | s3:ObjectAnnotation:* | s3:ObjectAnnotation:Put | s3:ObjectAnnotation:Delete | s3:ObjectRetention:Put

Required: No

Id

An optional unique identifier for configurations in a notification configuration. If you don't provide one, Amazon S3 will assign an ID.

Type: String

Required: No

Queue

The Amazon Resource Name (ARN) of the Amazon SQS queue to which Amazon S3 publishes a message when it detects events of the specified type.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

RecordExpiration

Service: Amazon S3

The journal table record expiration settings for a journal table in an S3 Metadata configuration.

Contents

Expiration

Specifies whether journal table record expiration is enabled or disabled.

Type: String

Valid Values: ENABLED | DISABLED

Required: Yes

Days

If you enable journal table record expiration, you can set the number of days to retain your journal table records. Journal table records must be retained for a minimum of 7 days. To set this value, specify any whole number from 7 to 2147483647. For example, to retain your journal table records for one year, set this value to 365.

Type: Integer

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

RecordsEvent

Service: Amazon S3

The container for the records event.

Contents

Payload

The byte array of partial, one or more result records. S3 Select doesn't guarantee that a record will be self-contained in one record frame. To ensure continuous streaming of data, S3 Select might split the same record across multiple record frames instead of aggregating the results in memory. Some S3 clients (for example, the AWS SDK for Java) handle this behavior by creating a `ByteArrayStream` out of the response by default. Other clients might not handle this behavior by default. In those cases, you must aggregate the results on the client side and parse the response.

Type: Base64-encoded binary data object

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Redirect

Service: Amazon S3

Specifies how requests are redirected. In the event of an error, you can specify a different error code to return.

Contents

HostName

The host name to use in the redirect request.

Type: String

Required: No

HttpRedirectCode

The HTTP redirect code to use on the response. Not required if one of the siblings is present.

Type: String

Required: No

Protocol

Protocol to use when redirecting requests. The default is the protocol that is used in the original request.

Type: String

Valid Values: http | https

Required: No

ReplaceKeyPrefixWith

The object key prefix to use in the redirect request. For example, to redirect requests for all pages with prefix docs/ (objects in the docs/ folder) to documents/, you can set a condition block with `KeyPrefixEquals` set to docs/ and in the Redirect set `ReplaceKeyPrefixWith` to /documents. Not required if one of the siblings is present. Can be present only if `ReplaceKeyWith` is not provided.

⚠ Important

Replacement must be made for object keys containing special characters (such as carriage returns) when using XML requests. For more information, see [XML related object key constraints](#).

Type: String

Required: No

ReplaceKeyWith

The specific object key to use in the redirect request. For example, redirect request to `error.html`. Not required if one of the siblings is present. Can be present only if `ReplaceKeyPrefixWith` is not provided.

⚠ Important

Replacement must be made for object keys containing special characters (such as carriage returns) when using XML requests. For more information, see [XML related object key constraints](#).

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

RedirectAllRequestsTo

Service: Amazon S3

Specifies the redirect behavior of all requests to a website endpoint of an Amazon S3 bucket.

Contents

HostName

Name of the host where requests are redirected.

Type: String

Required: Yes

Protocol

Protocol to use when redirecting requests. The default is the protocol that is used in the original request.

Type: String

Valid Values: http | https

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ReplicaModifications

Service: Amazon S3

A filter that you can specify for selection for modifications on replicas. Amazon S3 doesn't replicate replica modifications by default. In the latest version of replication configuration (when `Filter` is specified), you can specify this element and set the status to `Enabled` to replicate modifications on replicas.

Note

If you don't specify the `Filter` element, Amazon S3 assumes that the replication configuration is the earlier version, V1. In the earlier version, this element is not allowed.

Contents

Status

Specifies whether Amazon S3 replicates modifications on replicas.

Type: String

Valid Values: `Enabled` | `Disabled`

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ReplicationConfiguration

Service: Amazon S3

A container for replication rules. You can add up to 1,000 rules. The maximum size of a replication configuration is 2 MB.

Contents

Role

The Amazon Resource Name (ARN) of the AWS Identity and Access Management (IAM) role that Amazon S3 assumes when replicating objects. For more information, see [How to Set Up Replication](#) in the *Amazon S3 User Guide*.

Type: String

Required: Yes

Rules

A container for one or more replication rules. A replication configuration must have at least one rule and can contain a maximum of 1,000 rules.

Type: Array of [ReplicationRule](#) data types

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ReplicationRule

Service: Amazon S3

Specifies which Amazon S3 objects to replicate and where to store the replicas.

Contents

Destination

A container for information about the replication destination and its configurations including enabling the S3 Replication Time Control (S3 RTC).

Type: [Destination](#) data type

Required: Yes

Status

Specifies whether the rule is enabled.

Type: String

Valid Values: Enabled | Disabled

Required: Yes

DeleteMarkerReplication

Specifies whether Amazon S3 replicates delete markers. If you specify a `Filter` in your replication configuration, you must also include a `DeleteMarkerReplication` element. If your `Filter` includes a `Tag` element, the `DeleteMarkerReplication` Status must be set to `Disabled`, because Amazon S3 does not support replicating delete markers for tag-based rules. For an example configuration, see [Basic Rule Configuration](#).

For more information about delete marker replication, see [Basic Rule Configuration](#).

Note

If you are using an earlier version of the replication configuration, Amazon S3 handles replication of delete markers differently. For more information, see [Backward Compatibility](#).

Type: [DeleteMarkerReplication](#) data type

Required: No

ExistingObjectReplication

Optional configuration to replicate existing source bucket objects.

Note

This parameter is no longer supported. To replicate existing objects, see [Replicating existing objects with S3 Batch Replication](#) in the *Amazon S3 User Guide*.

Type: [ExistingObjectReplication](#) data type

Required: No

Filter

A filter that identifies the subset of objects to which the replication rule applies. A `Filter` must specify exactly one `Prefix`, `Tag`, or an `And` child element.

Type: [ReplicationRuleFilter](#) data type

Required: No

ID

A unique identifier for the rule. The maximum value is 255 characters.

Type: String

Required: No

Prefix

This member has been deprecated.

An object key name prefix that identifies the object or objects to which the rule applies. The maximum prefix length is 1,024 characters. To include all objects in a bucket, specify an empty string.

⚠ Important

Replacement must be made for object keys containing special characters (such as carriage returns) when using XML requests. For more information, see [XML related object key constraints](#).

Type: String

Required: No

Priority

The priority indicates which rule has precedence whenever two or more replication rules conflict. Amazon S3 will attempt to replicate objects according to all replication rules. However, if there are two or more rules with the same destination bucket, then objects will be replicated according to the rule with the highest priority. The higher the number, the higher the priority.

For more information, see [Replication](#) in the *Amazon S3 User Guide*.

Type: Integer

Required: No

SourceSelectionCriteria

A container that describes additional filters for identifying the source objects that you want to replicate. You can choose to enable or disable the replication of these objects. Currently, Amazon S3 supports only the filter that you can specify for objects created with server-side encryption using a customer managed key stored in AWS Key Management Service (SSE-KMS).

Type: [SourceSelectionCriteria](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ReplicationRuleAndOperator

Service: Amazon S3

A container for specifying rule filters. The filters determine the subset of objects to which the rule applies. This element is required only if you specify more than one filter.

For example:

- If you specify both a `Prefix` and a `Tag` filter, wrap these filters in an `And` tag.
- If you specify a filter based on multiple tags, wrap the `Tag` elements in an `And` tag.

Contents

Prefix

An object key name prefix that identifies the subset of objects to which the rule applies.

Type: String

Required: No

Tags

An array of tags containing key and value pairs.

Type: Array of [Tag](#) data types

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ReplicationRuleFilter

Service: Amazon S3

A filter that identifies the subset of objects to which the replication rule applies. A `Filter` must specify exactly one `Prefix`, `Tag`, or an `And` child element.

Contents

And

A container for specifying rule filters. The filters determine the subset of objects to which the rule applies. This element is required only if you specify more than one filter. For example:

- If you specify both a `Prefix` and a `Tag` filter, wrap these filters in an `And` tag.
- If you specify a filter based on multiple tags, wrap the `Tag` elements in an `And` tag.

Type: [ReplicationRuleAndOperator](#) data type

Required: No

Prefix

An object key name prefix that identifies the subset of objects to which the rule applies.

Important

Replacement must be made for object keys containing special characters (such as carriage returns) when using XML requests. For more information, see [XML related object key constraints](#).

Type: String

Required: No

Tag

A container for specifying a tag key and value.

The rule applies only to objects that have the tag in their tag set.

Type: [Tag](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ReplicationTime

Service: Amazon S3

A container specifying S3 Replication Time Control (S3 RTC) related information, including whether S3 RTC is enabled and the time when all objects and operations on objects must be replicated. Must be specified together with a `Metrics` block.

Contents

Status

Specifies whether the replication time is enabled.

Type: String

Valid Values: Enabled | Disabled

Required: Yes

Time

A container specifying the time by which replication should be complete for all objects and operations on objects.

Type: [ReplicationTimeValue](#) data type

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ReplicationTimeValue

Service: Amazon S3

A container specifying the time value for S3 Replication Time Control (S3 RTC) and replication metrics EventThreshold.

Contents

Minutes

Contains an integer specifying time in minutes.

Valid value: 15

Type: Integer

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

RequestPaymentConfiguration

Service: Amazon S3

Container for Payer.

Contents

Payer

Specifies who pays for the download and request fees.

Type: String

Valid Values: Requester | BucketOwner

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

RequestProgress

Service: Amazon S3

Container for specifying if periodic QueryProgress messages should be sent.

Contents

Enabled

Specifies whether periodic QueryProgress frames should be sent. Valid values: TRUE, FALSE.
Default value: FALSE.

Type: Boolean

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

RestoreRequest

Service: Amazon S3

Container for restore job parameters.

Contents

Days

Lifetime of the active copy in days. Do not use with restores that specify `OutputLocation`.

The `Days` element is required for regular restores, and must not be provided for select requests.

Type: Integer

Required: No

Description

The optional description for the job.

Type: String

Required: No

GlacierJobParameters

S3 Glacier related parameters pertaining to this job. Do not use with restores that specify `OutputLocation`.

Type: [GlacierJobParameters](#) data type

Required: No

OutputLocation

Describes the location where the restore job's output is stored.

Type: [OutputLocation](#) data type

Required: No

SelectParameters

Important

Amazon S3 Select is no longer available to new customers. Existing customers of Amazon S3 Select can continue to use the feature as usual. [Learn more](#)

Describes the parameters for Select job types.

Type: [SelectParameters](#) data type

Required: No

Tier

Retrieval tier at which the restore will be processed.

Type: String

Valid Values: Standard | Bulk | Expedited

Required: No

Type

Important

Amazon S3 Select is no longer available to new customers. Existing customers of Amazon S3 Select can continue to use the feature as usual. [Learn more](#)

Type of restore request.

Type: String

Valid Values: SELECT

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

RestoreStatus

Service: Amazon S3

Specifies the restoration status of an object. Objects in certain storage classes must be restored before they can be retrieved. For more information about these storage classes and how to work with archived objects, see [Working with archived objects](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported for directory buckets. Directory buckets only support EXPRESS_ONEZONE (the S3 Express One Zone storage class) in Availability Zones and ONEZONE_IA (the S3 One Zone-Infrequent Access storage class) in Dedicated Local Zones.

Contents

IsRestoreInProgress

Specifies whether the object is currently being restored. If the object restoration is in progress, the header returns the value TRUE. For example:

```
x-amz-optional-object-attributes: IsRestoreInProgress="true"
```

If the object restoration has completed, the header returns the value FALSE. For example:

```
x-amz-optional-object-attributes: IsRestoreInProgress="false",  
RestoreExpiryDate="2012-12-21T00:00:00.000Z"
```

If the object hasn't been restored, there is no header response.

Type: Boolean

Required: No

RestoreExpiryDate

Indicates when the restored copy will expire. This value is populated only if the object has already been restored. For example:

```
x-amz-optional-object-attributes: IsRestoreInProgress="false",  
RestoreExpiryDate="2012-12-21T00:00:00.000Z"
```

Type: Timestamp

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

RoutingRule

Service: Amazon S3

Specifies the redirect behavior and when a redirect is applied. For more information about routing rules, see [Configuring advanced conditional redirects](#) in the *Amazon S3 User Guide*.

Contents

Redirect

Container for redirect information. You can redirect requests to another host, to another page, or with another protocol. In the event of an error, you can specify a different error code to return.

Type: [Redirect](#) data type

Required: Yes

Condition

A container for describing a condition that must be met for the specified redirect to apply. For example, 1. If request is for pages in the /docs folder, redirect to the /documents folder. 2. If request results in HTTP error 4xx, redirect request to another host where you might process the error.

Type: [Condition](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Rule

Service: Amazon S3

Specifies lifecycle rules for an Amazon S3 bucket. For more information, see [Put Bucket Lifecycle Configuration](#) in the *Amazon S3 API Reference*. For examples, see [Put Bucket Lifecycle Configuration Examples](#).

Contents

Prefix

Object key prefix that identifies one or more objects to which this rule applies.

Important

Replacement must be made for object keys containing special characters (such as carriage returns) when using XML requests. For more information, see [XML related object key constraints](#).

Type: String

Required: Yes

Status

If Enabled, the rule is currently being applied. If Disabled, the rule is not currently being applied.

Type: String

Valid Values: Enabled | Disabled

Required: Yes

AbortIncompleteMultipartUpload

Specifies the days since the initiation of an incomplete multipart upload that Amazon S3 will wait before permanently removing all parts of the upload. For more information, see [Aborting Incomplete Multipart Uploads Using a Bucket Lifecycle Configuration](#) in the *Amazon S3 User Guide*.

Type: [AbortIncompleteMultipartUpload](#) data type

Required: No

Expiration

Specifies the expiration for the lifecycle of the object.

Type: [LifecycleExpiration](#) data type

Required: No

ID

Unique identifier for the rule. The value can't be longer than 255 characters.

Type: String

Required: No

NoncurrentVersionExpiration

Specifies when noncurrent object versions expire. Upon expiration, Amazon S3 permanently deletes the noncurrent object versions. You set this lifecycle configuration action on a bucket that has versioning enabled (or suspended) to request that Amazon S3 delete noncurrent object versions at a specific period in the object's lifetime.

Note

This parameter applies to general purpose buckets only. It is not supported for directory bucket lifecycle configurations.

Type: [NoncurrentVersionExpiration](#) data type

Required: No

NoncurrentVersionTransition

Container for the transition rule that describes when noncurrent objects transition to the STANDARD_IA, ONEZONE_IA, INTELLIGENT_TIERING, GLACIER_IR, GLACIER, or DEEP_ARCHIVE storage class. If your bucket is versioning-enabled (or versioning is suspended), you can set this action to request that Amazon S3 transition noncurrent object versions to

the STANDARD_IA, ONEZONE_IA, INTELLIGENT_TIERING, GLACIER_IR, GLACIER, or DEEP_ARCHIVE storage class at a specific period in the object's lifetime.

Type: [NoncurrentVersionTransition](#) data type

Required: No

Transition

Specifies when an object transitions to a specified storage class. For more information about Amazon S3 lifecycle configuration rules, see [Transitioning Objects Using Amazon S3 Lifecycle](#) in the *Amazon S3 User Guide*.

Type: [Transition](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3KeyFilter

Service: Amazon S3

A container for object key name prefix and suffix filtering rules.

Contents

FilterRules

A list of containers for the key-value pair that defines the criteria for the filter rule.

Type: Array of [FilterRule](#) data types

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3Location

Service: Amazon S3

Describes an Amazon S3 location that will receive the results of the restore request.

Contents

BucketName

The name of the bucket where the restore results will be placed.

Type: String

Required: Yes

Prefix

The prefix that is prepended to the restore results for this request.

Type: String

Required: Yes

AccessControlList

A list of grants that control access to the staged results.

Type: Array of [Grant](#) data types

Required: No

CannedACL

The canned ACL to apply to the restore results.

Type: String

Valid Values: `private` | `public-read` | `public-read-write` | `authenticated-read` | `aws-exec-read` | `bucket-owner-read` | `bucket-owner-full-control`

Required: No

Encryption

Contains the type of server-side encryption used.

Type: [Encryption](#) data type

Required: No

StorageClass

The class of storage used to store the restore results.

Type: String

Valid Values: STANDARD | REDUCED_REDUNDANCY | STANDARD_IA | ONEZONE_IA | INTELLIGENT_TIERING | GLACIER | DEEP_ARCHIVE | OUTPOSTS | GLACIER_IR | SNOW | EXPRESS_ONEZONE | FSX_OPENZFS | FSX_ONTAP | AWS_BACKUP_WARM | AWS_BACKUP_LOW_COST_WARM

Required: No

Tagging

The tag-set that is applied to the restore results.

Type: [Tagging](#) data type

Required: No

UserMetadata

A list of metadata to store with the restore results in S3.

Type: Array of [MetadataEntry](#) data types

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3TablesDestination

Service: Amazon S3

The destination information for a V1 S3 Metadata configuration. The destination table bucket must be in the same Region and AWS account as the general purpose bucket. The specified metadata table name must be unique within the `aws_s3_metadata` namespace in the destination table bucket.

Note

If you created your S3 Metadata configuration before July 15, 2025, we recommend that you delete and re-create your configuration by using [CreateBucketMetadataConfiguration](#) so that you can expire journal table records and create a live inventory table.

Contents

TableBucketArn

The Amazon Resource Name (ARN) for the table bucket that's specified as the destination in the metadata table configuration. The destination table bucket must be in the same Region and AWS account as the general purpose bucket.

Type: String

Required: Yes

TableName

The name for the metadata table in your metadata table configuration. The specified metadata table name must be unique within the `aws_s3_metadata` namespace in the destination table bucket.

Type: String

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3TablesDestinationResult

Service: Amazon S3

The destination information for a V1 S3 Metadata configuration. The destination table bucket must be in the same Region and AWS account as the general purpose bucket. The specified metadata table name must be unique within the `aws_s3_metadata` namespace in the destination table bucket.

Note

If you created your S3 Metadata configuration before July 15, 2025, we recommend that you delete and re-create your configuration by using [CreateBucketMetadataConfiguration](#) so that you can expire journal table records and create a live inventory table.

Contents

TableArn

The Amazon Resource Name (ARN) for the metadata table in the metadata table configuration. The specified metadata table name must be unique within the `aws_s3_metadata` namespace in the destination table bucket.

Type: String

Required: Yes

TableBucketArn

The Amazon Resource Name (ARN) for the table bucket that's specified as the destination in the metadata table configuration. The destination table bucket must be in the same Region and AWS account as the general purpose bucket.

Type: String

Required: Yes

TableName

The name for the metadata table in your metadata table configuration. The specified metadata table name must be unique within the `aws_s3_metadata` namespace in the destination table bucket.

Type: String

Required: Yes

TableNameNamespace

The table bucket namespace for the metadata table in your metadata table configuration. This value is always `aws_s3_metadata`.

Type: String

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ScanRange

Service: Amazon S3

Specifies the byte range of the object to get the records from. A record is processed when its first byte is contained by the range. This parameter is optional, but when specified, it must not be empty. See RFC 2616, Section 14.35.1 about how to specify the start and end of the range.

Contents

End

Specifies the end of the byte range. This parameter is optional. Valid values: non-negative integers. The default value is one less than the size of the object being queried. If only the End parameter is supplied, it is interpreted to mean scan the last N bytes of the file. For example, `<scanrange><end>50</end></scanrange>` means scan the last 50 bytes.

Type: Long

Required: No

Start

Specifies the start of the byte range. This parameter is optional. Valid values: non-negative integers. The default value is 0. If only start is supplied, it means scan from that point to the end of the file. For example, `<scanrange><start>50</start></scanrange>` means scan from byte 50 until the end of the file.

Type: Long

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SelectObjectContentEventStream

Service: Amazon S3

The container for selecting objects from a content event stream.

Contents

Cont

The Continuation Event.

Type: [ContinuationEvent](#) data type

Required: No

End

The End Event.

Type: [EndEvent](#) data type

Required: No

Progress

The Progress Event.

Type: [ProgressEvent](#) data type

Required: No

Records

The Records Event.

Type: [RecordsEvent](#) data type

Required: No

Stats

The Stats Event.

Type: [StatsEvent](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SelectParameters

Service: Amazon S3

Important

Amazon S3 Select is no longer available to new customers. Existing customers of Amazon S3 Select can continue to use the feature as usual. [Learn more](#)

Describes the parameters for Select job types.

Learn [How to optimize querying your data in Amazon S3](#) using [Amazon Athena](#), [S3 Object Lambda](#), or client-side filtering.

Contents

Expression

Important

Amazon S3 Select is no longer available to new customers. Existing customers of Amazon S3 Select can continue to use the feature as usual. [Learn more](#)

The expression that is used to query the object.

Type: String

Required: Yes

ExpressionType

The type of the provided expression (for example, SQL).

Type: String

Valid Values: SQL

Required: Yes

InputSerialization

Describes the serialization format of the object.

Type: [InputSerialization](#) data type

Required: Yes

OutputSerialization

Describes how the results of the Select job are serialized.

Type: [OutputSerialization](#) data type

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ServerSideEncryptionByDefault

Service: Amazon S3

Describes the default server-side encryption to apply to new objects in the bucket. If a PUT Object request doesn't specify any server-side encryption, this default encryption will be applied. For more information, see [PutBucketEncryption](#).

Note

- **General purpose buckets** - If you don't specify a customer managed key at configuration, Amazon S3 automatically creates an AWS KMS key (aws/s3) in your AWS account the first time that you add an object encrypted with SSE-KMS to a bucket. By default, Amazon S3 uses this KMS key for SSE-KMS.
- **Directory buckets** - Your SSE-KMS configuration can only support 1 [customer managed key](#) per directory bucket's lifetime. The [AWS managed key](#) (aws/s3) isn't supported.
- **Directory buckets** - For directory buckets, there are only two supported options for server-side encryption: SSE-S3 and SSE-KMS.

Contents

SSEAlgorithm

Server-side encryption algorithm to use for the default encryption.

Note

For directory buckets, there are only two supported values for server-side encryption: AES256 and aws:kms.

Type: String

Valid Values: AES256 | aws:fsx | aws:backup | aws:kms | aws:kms:dsse

Required: Yes

KMSMasterKeyID

AWS Key Management Service (KMS) customer managed key ID to use for the default encryption.

Note

- **General purpose buckets** - This parameter is allowed if and only if `SSEAlgorithm` is set to `aws:kms` or `aws:kms:dsse`.
- **Directory buckets** - This parameter is allowed if and only if `SSEAlgorithm` is set to `aws:kms`.

You can specify the key ID, key alias, or the Amazon Resource Name (ARN) of the KMS key.

- Key ID: `1234abcd-12ab-34cd-56ef-1234567890ab`
- Key ARN: `arn:aws:kms:us-east-2:111122223333:key/1234abcd-12ab-34cd-56ef-1234567890ab`
- Key Alias: `alias/alias-name`

If you are using encryption with cross-account or AWS service operations, you must use a fully qualified KMS key ARN. For more information, see [Using encryption for cross-account operations](#).

Note

- **General purpose buckets** - If you're specifying a customer managed KMS key, we recommend using a fully qualified KMS key ARN. If you use a KMS key alias instead, then AWS KMS resolves the key within the requester's account. This behavior can result in data that's encrypted with a KMS key that belongs to the requester, and not the bucket owner. Also, if you use a key ID, you can run into a `LogDestination undeliverable` error when creating a VPC flow log.
- **Directory buckets** - When you specify an [AWS KMS customer managed key](#) for encryption in your directory bucket, only use the key ID or key ARN. The key alias format of the KMS key isn't supported.

** Important**

Amazon S3 only supports symmetric encryption KMS keys. For more information, see [Asymmetric keys in AWS KMS](#) in the *AWS Key Management Service Developer Guide*.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ServerSideEncryptionConfiguration

Service: Amazon S3

Specifies the default server-side-encryption configuration.

Contents

Rules

Container for information about a particular server-side encryption configuration rule.

Type: Array of [ServerSideEncryptionRule](#) data types

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ServerSideEncryptionRule

Service: Amazon S3

Specifies the default server-side encryption configuration.

Note

- **General purpose buckets** - If you're specifying a customer managed KMS key, we recommend using a fully qualified KMS key ARN. If you use a KMS key alias instead, then AWS KMS resolves the key within the requester's account. This behavior can result in data that's encrypted with a KMS key that belongs to the requester, and not the bucket owner.
- **Directory buckets** - When you specify an [AWS KMS customer managed key](#) for encryption in your directory bucket, only use the key ID or key ARN. The key alias format of the KMS key isn't supported.

Contents

ApplyServerSideEncryptionByDefault

Specifies the default server-side encryption to apply to new objects in the bucket. If a PUT Object request doesn't specify any server-side encryption, this default encryption will be applied.

Type: [ServerSideEncryptionByDefault](#) data type

Required: No

BlockedEncryptionTypes

A bucket-level setting for Amazon S3 general purpose buckets used to prevent the upload of new objects encrypted with the specified server-side encryption type. For example, blocking an encryption type will block PutObject, CopyObject, PostObject, multipart upload, and replication requests to the bucket for objects with the specified encryption type. However, you can continue to read and list any pre-existing objects already encrypted with the specified encryption type. For more information, see [Blocking or unblocking SSE-C for a general purpose bucket](#).

Note

Currently, this parameter only supports blocking or unblocking server-side encryption with customer-provided keys (SSE-C). For more information about SSE-C, see [Using server-side encryption with customer-provided keys \(SSE-C\)](#).

Type: [BlockedEncryptionTypes](#) data type

Required: No

BucketKeyEnabled

Specifies whether Amazon S3 should use an S3 Bucket Key with server-side encryption using KMS (SSE-KMS) for new objects in the bucket. Existing objects are not affected. Setting the `BucketKeyEnabled` element to `true` causes Amazon S3 to use an S3 Bucket Key.

Note

- **General purpose buckets** - By default, S3 Bucket Key is not enabled. For more information, see [Amazon S3 Bucket Keys](#) in the *Amazon S3 User Guide*.
- **Directory buckets** - S3 Bucket Keys are always enabled for GET and PUT operations in a directory bucket and can't be disabled. S3 Bucket Keys aren't supported, when you copy SSE-KMS encrypted objects from general purpose buckets to directory buckets, from directory buckets to general purpose buckets, or between directory buckets, through [CopyObject](#), [UploadPartCopy](#), [the Copy operation in Batch Operations](#), or [the import jobs](#). In this case, Amazon S3 makes a call to AWS KMS every time a copy request is made for a KMS-encrypted object.

Type: Boolean

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SessionCredentials

Service: Amazon S3

The established temporary security credentials of the session.

Note

Directory buckets - These session credentials are only supported for the authentication and authorization of Zonal endpoint API operations on directory buckets.

Contents

AccessKeyId

A unique identifier that's associated with a secret access key. The access key ID and the secret access key are used together to sign programmatic AWS requests cryptographically.

Type: String

Required: Yes

Expiration

Temporary security credentials expire after a specified interval. After temporary credentials expire, any calls that you make with those credentials will fail. So you must generate a new set of temporary credentials. Temporary credentials cannot be extended or refreshed beyond the original specified interval.

Type: Timestamp

Required: Yes

SecretAccessKey

A key that's used with the access key ID to cryptographically sign programmatic AWS requests. Signing a request identifies the sender and prevents the request from being altered.

Type: String

Required: Yes

SessionToken

A part of the temporary security credentials. The session token is used to validate the temporary security credentials.

Type: String

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SimplePrefix

Service: Amazon S3

To use simple format for S3 keys for log objects, set SimplePrefix to an empty object.

[DestinationPrefix][YYYY]-[MM]-[DD]-[hh]-[mm]-[ss]-[UniqueString]

Contents

The members of this exception structure are context-dependent.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SourceSelectionCriteria

Service: Amazon S3

A container that describes additional filters for identifying the source objects that you want to replicate. You can choose to enable or disable the replication of these objects. Currently, Amazon S3 supports only the filter that you can specify for objects created with server-side encryption using a customer managed key stored in AWS Key Management Service (SSE-KMS).

Contents

ReplicaModifications

A filter that you can specify for selections for modifications on replicas. Amazon S3 doesn't replicate replica modifications by default. In the latest version of replication configuration (when `Filter` is specified), you can specify this element and set the status to `Enabled` to replicate modifications on replicas.

Note

If you don't specify the `Filter` element, Amazon S3 assumes that the replication configuration is the earlier version, V1. In the earlier version, this element is not allowed

Type: [ReplicaModifications](#) data type

Required: No

SseKmsEncryptedObjects

A container for filter information for the selection of Amazon S3 objects encrypted with AWS KMS. If you include `SourceSelectionCriteria` in the replication configuration, this element is required.

Type: [SseKmsEncryptedObjects](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SSEKMS

Service: Amazon S3

Specifies the use of SSE-KMS to encrypt delivered inventory reports.

Contents

KeyId

Specifies the ID of the AWS Key Management Service (AWS KMS) symmetric encryption customer managed key to use for encrypting inventory reports.

Type: String

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SseKmsEncryptedObjects

Service: Amazon S3

A container for filter information for the selection of S3 objects encrypted with AWS KMS.

Contents

Status

Specifies whether Amazon S3 replicates objects created with server-side encryption using an AWS KMS key stored in AWS Key Management Service.

Type: String

Valid Values: Enabled | Disabled

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SSEKMSEncryption

Service: Amazon S3

If SSEKMS is specified for `ObjectEncryption`, this data type specifies the AWS KMS key Amazon Resource Name (ARN) to use and whether to use an S3 Bucket Key for server-side encryption using AWS Key Management Service (AWS KMS) keys (SSE-KMS).

Contents

KMSKeyArn

Specifies the AWS KMS key Amazon Resource Name (ARN) to use for the updated server-side encryption type. Required if `ObjectEncryption` specifies SSEKMS.

Note

You must specify the full AWS KMS key ARN. The KMS key ID and KMS key alias aren't supported.

Pattern: `(arn:aws[-a-z0-9]*:kms:[-a-z0-9]*:[0-9]{12}:key/.+)`

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `arn:aws[a-zA-Z0-9-]*:kms:[a-z0-9-]+:[0-9]{12}:key/[a-zA-Z0-9-]+`

Required: Yes

BucketKeyEnabled

Specifies whether Amazon S3 should use an S3 Bucket Key for object encryption with server-side encryption using AWS Key Management Service (AWS KMS) keys (SSE-KMS). If this value isn't specified, it defaults to `false`. Setting this value to `true` causes Amazon S3 to use an S3 Bucket Key for object encryption with SSE-KMS. For more information, see [Using Amazon S3 Bucket Keys](#) in the *Amazon S3 User Guide*.

Valid Values: `true` | `false`

Type: Boolean

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SSES3

Service: Amazon S3

Specifies the use of SSE-S3 to encrypt delivered inventory reports.

Contents

The members of this exception structure are context-dependent.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Stats

Service: Amazon S3

Container for the stats details.

Contents

BytesProcessed

The total number of uncompressed object bytes processed.

Type: Long

Required: No

BytesReturned

The total number of bytes of records payload data returned.

Type: Long

Required: No

BytesScanned

The total number of object bytes scanned.

Type: Long

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

StatsEvent

Service: Amazon S3

Container for the Stats Event.

Contents

Details

The Stats event details.

Type: [Stats](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

StorageClassAnalysis

Service: Amazon S3

Specifies data related to access patterns to be collected and made available to analyze the tradeoffs between different storage classes for an Amazon S3 bucket.

Contents

DataExport

Specifies how data related to the storage class analysis for an Amazon S3 bucket should be exported.

Type: [StorageClassAnalysisDataExport](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

StorageClassAnalysisDataExport

Service: Amazon S3

Container for data related to the storage class analysis for an Amazon S3 bucket for export.

Contents

Destination

The place to store the data for an analysis.

Type: [AnalyticsExportDestination](#) data type

Required: Yes

OutputSchemaVersion

The version of the output schema to use when exporting data. Must be V_1.

Type: String

Valid Values: V_1

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Tag

Service: Amazon S3

A container of a key value name pair.

Contents

Key

Name of the object key.

Type: String

Length Constraints: Minimum length of 1.

Required: Yes

Value

Value of the tag.

Type: String

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Tagging

Service: Amazon S3

Container for TagSet elements.

Contents

TagSet

A collection for a set of tags

Type: Array of [Tag](#) data types

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

TargetGrant

Service: Amazon S3

Container for granting information.

Buckets that use the bucket owner enforced setting for Object Ownership don't support target grants. For more information, see [Permissions server access log delivery](#) in the *Amazon S3 User Guide*.

Contents

Grantee

Container for the person being granted permissions.

Type: [Grantee](#) data type

Required: No

Permission

Logging permissions assigned to the grantee for the bucket.

Type: String

Valid Values: FULL_CONTROL | READ | WRITE

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

TargetObjectKeyFormat

Service: Amazon S3

Amazon S3 key format for log objects. Only one format, PartitionedPrefix or SimplePrefix, is allowed.

Contents

PartitionedPrefix

Partitioned S3 key for log objects.

Type: [PartitionedPrefix](#) data type

Required: No

SimplePrefix

To use the simple format for S3 keys for log objects. To specify SimplePrefix format, set SimplePrefix to {}.

Type: [SimplePrefix](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Tiering

Service: Amazon S3

The S3 Intelligent-Tiering storage class is designed to optimize storage costs by automatically moving data to the most cost-effective storage access tier, without additional operational overhead.

Contents

AccessTier

S3 Intelligent-Tiering access tier. See [Storage class for automatically optimizing frequently and infrequently accessed objects](#) for a list of access tiers in the S3 Intelligent-Tiering storage class.

Type: String

Valid Values: ARCHIVE_ACCESS | DEEP_ARCHIVE_ACCESS

Required: Yes

Days

The number of consecutive days of no access after which an object will be eligible to be transitioned to the corresponding tier. The minimum number of days specified for Archive Access tier must be at least 90 days and Deep Archive Access tier must be at least 180 days. The maximum can be up to 2 years (730 days).

Type: Integer

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

TopicConfiguration

Service: Amazon S3

A container for specifying the configuration for publication of messages to an Amazon Simple Notification Service (Amazon SNS) topic when Amazon S3 detects specified events.

Contents

Events

The Amazon S3 bucket event about which to send notifications. For more information, see [Supported Event Types](#) in the *Amazon S3 User Guide*.

Type: Array of strings

Valid Values: s3:ReducedRedundancyLostObject | s3:ObjectCreated:* | s3:ObjectCreated:Put | s3:ObjectCreated:Post | s3:ObjectCreated:Copy | s3:ObjectCreated:CompleteMultipartUpload | s3:ObjectRemoved:* | s3:ObjectRemoved:Delete | s3:ObjectRemoved:DeleteMarkerCreated | s3:ObjectRestore:* | s3:ObjectRestore:Post | s3:ObjectRestore:Completed | s3:Replication:* | s3:Replication:OperationFailedReplication | s3:Replication:OperationNotTracked | s3:Replication:OperationMissedThreshold | s3:Replication:OperationReplicatedAfterThreshold | s3:ObjectRestore:Delete | s3:LifecycleTransition | s3:IntelligentTiering | s3:ObjectAcl:Put | s3:LifecycleExpiration:* | s3:LifecycleExpiration:Delete | s3:LifecycleExpiration:DeleteMarkerCreated | s3:ObjectTagging:* | s3:ObjectTagging:Put | s3:ObjectTagging:Delete | s3:ObjectAnnotation:* | s3:ObjectAnnotation:Put | s3:ObjectAnnotation:Delete | s3:ObjectRetention:Put

Required: Yes

TopicArn

The Amazon Resource Name (ARN) of the Amazon SNS topic to which Amazon S3 publishes a message when it detects events of the specified type.

Type: String

Required: Yes

Filter

Specifies object key name filtering rules. For information about key name filtering, see [Configuring event notifications using object key name filtering](#) in the *Amazon S3 User Guide*.

Type: [NotificationConfigurationFilter](#) data type

Required: No

Id

An optional unique identifier for configurations in a notification configuration. If you don't provide one, Amazon S3 will assign an ID.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

TopicConfigurationDeprecated

Service: Amazon S3

A container for specifying the configuration for publication of messages to an Amazon Simple Notification Service (Amazon SNS) topic when Amazon S3 detects specified events. This data type is deprecated. Use [TopicConfiguration](#) instead.

Contents

Event

This member has been deprecated.

Bucket event for which to send notifications.

Type: String

Valid Values: s3:ReducedRedundancyLostObject | s3:ObjectCreated:* | s3:ObjectCreated:Put | s3:ObjectCreated:Post | s3:ObjectCreated:Copy | s3:ObjectCreated:CompleteMultipartUpload | s3:ObjectRemoved:* | s3:ObjectRemoved:Delete | s3:ObjectRemoved:DeleteMarkerCreated | s3:ObjectRestore:* | s3:ObjectRestore:Post | s3:ObjectRestore:Completed | s3:Replication:* | s3:Replication:OperationFailedReplication | s3:Replication:OperationNotTracked | s3:Replication:OperationMissedThreshold | s3:Replication:OperationReplicatedAfterThreshold | s3:ObjectRestore:Delete | s3:LifecycleTransition | s3: IntelligentTiering | s3:ObjectAcl:Put | s3:LifecycleExpiration:* | s3:LifecycleExpiration:Delete | s3:LifecycleExpiration:DeleteMarkerCreated | s3:ObjectTagging:* | s3:ObjectTagging:Put | s3:ObjectTagging:Delete | s3:ObjectAnnotation:* | s3:ObjectAnnotation:Put | s3:ObjectAnnotation:Delete | s3:ObjectRetention:Put

Required: No

Events

A collection of events related to objects

Type: Array of strings

Valid Values: s3:ReducedRedundancyLostObject | s3:ObjectCreated:* | s3:ObjectCreated:Put | s3:ObjectCreated:Post | s3:ObjectCreated:Copy | s3:ObjectCreated:CompleteMultipartUpload | s3:ObjectRemoved:* | s3:ObjectRemoved:Delete | s3:ObjectRemoved:DeleteMarkerCreated | s3:ObjectRestore:* | s3:ObjectRestore:Post | s3:ObjectRestore:Completed | s3:Replication:* | s3:Replication:OperationFailedReplication | s3:Replication:OperationNotTracked | s3:Replication:OperationMissedThreshold | s3:Replication:OperationReplicatedAfterThreshold | s3:ObjectRestore:Delete | s3:LifecycleTransition | s3:IntelligentTiering | s3:ObjectAcl:Put | s3:LifecycleExpiration:* | s3:LifecycleExpiration:Delete | s3:LifecycleExpiration:DeleteMarkerCreated | s3:ObjectTagging:* | s3:ObjectTagging:Put | s3:ObjectTagging:Delete | s3:ObjectAnnotation:* | s3:ObjectAnnotation:Put | s3:ObjectAnnotation:Delete | s3:ObjectRetention:Put

Required: No

Id

An optional unique identifier for configurations in a notification configuration. If you don't provide one, Amazon S3 will assign an ID.

Type: String

Required: No

Topic

Amazon SNS topic to which Amazon S3 will publish a message to report the specified events for the bucket.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Transition

Service: Amazon S3

Specifies when an object transitions to a specified storage class. For more information about Amazon S3 lifecycle configuration rules, see [Transitioning Objects Using Amazon S3 Lifecycle](#) in the *Amazon S3 User Guide*.

Contents

Date

Indicates when objects are transitioned to the specified storage class. The date value must be in ISO 8601 format. The time is always midnight UTC.

Type: Timestamp

Required: No

Days

Indicates the number of days after creation when objects are transitioned to the specified storage class. The value can be 0 or any positive integer. Be aware that some storage classes have a minimum storage duration and that you're charged for transitioning objects before their minimum storage duration. For more information, see [Constraints and considerations for transitions](#) in the *Amazon S3 User Guide*.

Type: Integer

Required: No

StorageClass

The storage class to which you want the object to transition.

Type: String

Valid Values: GLACIER | STANDARD_IA | ONEZONE_IA | INTELLIGENT_TIERING | DEEP_ARCHIVE | GLACIER_IR

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

VersioningConfiguration

Service: Amazon S3

Describes the versioning state of an Amazon S3 bucket. For more information, see [PUT Bucket versioning](#) in the *Amazon S3 API Reference*.

Contents

MFADelete

Specifies whether MFA delete is enabled in the bucket versioning configuration. This element is only returned if the bucket has been configured with MFA delete. If the bucket has never been so configured, this element is not returned.

Type: String

Valid Values: Enabled | Disabled

Required: No

Status

The versioning state of the bucket.

Type: String

Valid Values: Enabled | Suspended

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

WebsiteConfiguration

Service: Amazon S3

Specifies website configuration parameters for an Amazon S3 bucket.

Contents

ErrorDocument

The name of the error document for the website.

Type: [ErrorDocument](#) data type

Required: No

IndexDocument

The name of the index document for the website.

Type: [IndexDocument](#) data type

Required: No

RedirectAllRequestsTo

The redirect behavior for every request to this bucket's website endpoint.

Important

If you specify this property, you can't specify any other property.

Type: [RedirectAllRequestsTo](#) data type

Required: No

RoutingRules

Rules that define when a redirect is applied and the redirect behavior.

Type: Array of [RoutingRule](#) data types

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Amazon S3 Control

The following data types are supported by Amazon S3 Control:

- [AbortIncompleteMultipartUpload](#)
- [AccessControlTranslation](#)
- [AccessGrantsLocationConfiguration](#)
- [AccessPoint](#)
- [AccountLevel](#)
- [ActivityMetrics](#)
- [AdvancedCostOptimizationMetrics](#)
- [AdvancedDataProtectionMetrics](#)
- [AdvancedPerformanceMetrics](#)
- [AsyncErrorDetails](#)
- [AsyncOperation](#)
- [AsyncRequestParameters](#)
- [AsyncResponseDetails](#)
- [AwsLambdaTransformation](#)
- [BucketLevel](#)
- [CloudWatchMetrics](#)
- [CreateBucketConfiguration](#)
- [CreateMultiRegionAccessPointInput](#)
- [Credentials](#)
- [DeleteMarkerReplication](#)

- [DeleteMultiRegionAccessPointInput](#)
- [Destination](#)
- [DetailedStatusCodesMetrics](#)
- [DSSEKMSFilter](#)
- [EncryptionConfiguration](#)
- [EstablishedMultiRegionAccessPointPolicy](#)
- [Exclude](#)
- [ExistingObjectReplication](#)
- [GeneratedManifestEncryption](#)
- [Grantee](#)
- [Include](#)
- [JobDescriptor](#)
- [JobFailure](#)
- [JobListDescriptor](#)
- [JobManifest](#)
- [JobManifestGenerator](#)
- [JobManifestGeneratorFilter](#)
- [JobManifestLocation](#)
- [JobManifestSpec](#)
- [JobOperation](#)
- [JobProgressSummary](#)
- [JobReport](#)
- [JobTimers](#)
- [KeyNameConstraint](#)
- [LambdaInvokeOperation](#)
- [LifecycleConfiguration](#)
- [LifecycleExpiration](#)
- [LifecycleRule](#)
- [LifecycleRuleAndOperator](#)
- [LifecycleRuleFilter](#)

- [ListAccessGrantEntry](#)
- [ListAccessGrantsInstanceEntry](#)
- [ListAccessGrantsLocationsEntry](#)
- [ListCallerAccessGrantsEntry](#)
- [ListStorageLensConfigurationEntry](#)
- [ListStorageLensGroupEntry](#)
- [MatchObjectAge](#)
- [MatchObjectSize](#)
- [Metrics](#)
- [MultiRegionAccessPointPolicyDocument](#)
- [MultiRegionAccessPointRegionalResponse](#)
- [MultiRegionAccessPointReport](#)
- [MultiRegionAccessPointRoute](#)
- [MultiRegionAccessPointsAsyncResponse](#)
- [NoncurrentVersionExpiration](#)
- [NoncurrentVersionTransition](#)
- [NotSSEFilter](#)
- [ObjectEncryption](#)
- [ObjectEncryptionFilter](#)
- [ObjectLambdaAccessPoint](#)
- [ObjectLambdaAccessPointAlias](#)
- [ObjectLambdaConfiguration](#)
- [ObjectLambdaContentTransformation](#)
- [ObjectLambdaTransformationConfiguration](#)
- [PolicyStatus](#)
- [PrefixLevel](#)
- [PrefixLevelStorageMetrics](#)
- [ProposedMultiRegionAccessPointPolicy](#)
- [PublicAccessBlockConfiguration](#)
- [PutMultiRegionAccessPointPolicyInput](#)

- [Region](#)
- [RegionalBucket](#)
- [RegionReport](#)
- [ReplicaModifications](#)
- [ReplicationConfiguration](#)
- [ReplicationRule](#)
- [ReplicationRuleAndOperator](#)
- [ReplicationRuleFilter](#)
- [ReplicationTime](#)
- [ReplicationTimeValue](#)
- [S3AccessControlList](#)
- [S3AccessControlPolicy](#)
- [S3BucketDestination](#)
- [S3ComputeObjectChecksumOperation](#)
- [S3CopyObjectOperation](#)
- [S3DeleteObjectTaggingOperation](#)
- [S3GeneratedManifestDescriptor](#)
- [S3Grant](#)
- [S3Grantee](#)
- [S3InitiateRestoreObjectOperation](#)
- [S3JobManifestGenerator](#)
- [S3ManifestOutputLocation](#)
- [S3ObjectLockEventHoldDuration](#)
- [S3ObjectLockLegalHold](#)
- [S3ObjectLockRetentionEventHoldDuration](#)
- [S3ObjectMetadata](#)
- [S3ObjectOwner](#)
- [S3ReplicateObjectOperation](#)
- [S3Retention](#)
- [S3SetObjectAclOperation](#)

- [S3SetObjectLegalHoldOperation](#)
- [S3SetObjectRetentionOperation](#)
- [S3SetObjectTaggingOperation](#)
- [S3Tag](#)
- [S3UpdateObjectEncryptionOperation](#)
- [S3UpdateObjectEncryptionSSEKMS](#)
- [Scope](#)
- [SelectionCriteria](#)
- [SourceSelectionCriteria](#)
- [SSECFilter](#)
- [SSEKMS](#)
- [SseKmsEncryptedObjects](#)
- [SSEKMSEncryption](#)
- [SSEKMSTFilter](#)
- [SSES3](#)
- [SSES3Encryption](#)
- [SSES3Filter](#)
- [StorageLensAwsOrg](#)
- [StorageLensConfiguration](#)
- [StorageLensDataExport](#)
- [StorageLensDataExportEncryption](#)
- [StorageLensExpandedPrefixesDataExport](#)
- [StorageLensGroup](#)
- [StorageLensGroupAndOperator](#)
- [StorageLensGroupFilter](#)
- [StorageLensGroupLevel](#)
- [StorageLensGroupLevelSelectionCriteria](#)
- [StorageLensGroupOrOperator](#)
- [StorageLensTableDestination](#)
- [StorageLensTag](#)

- [Tag](#)
- [Tagging](#)
- [Transition](#)
- [VersioningConfiguration](#)
- [VpcConfiguration](#)

AbortIncompleteMultipartUpload

Service: Amazon S3 Control

The container for abort incomplete multipart upload

Contents

DaysAfterInitiation

Specifies the number of days after which Amazon S3 aborts an incomplete multipart upload to the Outposts bucket.

Type: Integer

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

AccessControlTranslation

Service: Amazon S3 Control

A container for information about access control for replicas.

Note

This is not supported by Amazon S3 on Outposts buckets.

Contents

Owner

Specifies the replica ownership.

Type: String

Valid Values: Destination

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

AccessGrantsLocationConfiguration

Service: Amazon S3 Control

The configuration options of the S3 Access Grants location. It contains the `S3SubPrefix` field. The grant scope, the data to which you are granting access, is the result of appending the `Subprefix` field to the scope of the registered location.

Contents

S3SubPrefix

The `S3SubPrefix` is appended to the location scope creating the grant scope. Use this field to narrow the scope of the grant to a subset of the location scope. This field is required if the location scope is the default location `s3://` because you cannot create a grant for all of your S3 data in the Region and must narrow the scope. For example, if the location scope is the default location `s3://`, the `S3SubPrefix` can be a `<bucket-name>/*`, so the full grant scope path would be `s3://<bucket-name>/*`. Or the `S3SubPrefix` can be `<bucket-name>/<prefix-name>*`, so the full grant scope path would be or `s3://<bucket-name>/<prefix-name>*`.

If the `S3SubPrefix` includes a prefix, append the wildcard character `*` after the prefix to indicate that you want to include all object key names in the bucket that start with that prefix.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2000.

Pattern: `^.+`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

AccessPoint

Service: Amazon S3 Control

An access point used to access a bucket.

Contents

Bucket

The name of the bucket associated with this access point.

Type: String

Length Constraints: Maximum length of 255.

Required: Yes

Name

The name of this access point.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

NetworkOrigin

Indicates whether this access point allows access from the public internet. If `VpcConfiguration` is specified for this access point, then `NetworkOrigin` is `VPC`, and the access point doesn't allow access from the public internet. Otherwise, `NetworkOrigin` is `Internet`, and the access point allows access from the public internet, subject to the access point and bucket access policies.

Type: String

Valid Values: `Internet` | `VPC`

Required: Yes

AccessPointArn

The ARN for the access point.

Type: String

Length Constraints: Minimum length of 4. Maximum length of 128.

Required: No

Alias

The name or alias of the access point.

Type: String

Length Constraints: Maximum length of 63.

Pattern: `^[0-9a-z\\-]{63}`

Required: No

BucketAccountId

The AWS account ID associated with the S3 bucket associated with this access point.

Type: String

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: No

DataSourceId

A unique identifier for the data source of the access point.

Type: String

Length Constraints: Maximum length of 191.

Required: No

DataSourceType

The type of the data source that the access point is attached to.

Type: String

Required: No

VpcConfiguration

The virtual private cloud (VPC) configuration for this access point, if one exists.

Note

This element is empty if this access point is an Amazon S3 on Outposts access point that is used by other AWS services.

Type: [VpcConfiguration](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

AccountLevel

Service: Amazon S3 Control

A container element for the account-level Amazon S3 Storage Lens configuration.

Note

You must enable Storage Lens metrics consistently at both the account level and bucket level, or your request will fail.

For more information about S3 Storage Lens, see [Assessing your storage activity and usage with S3 Storage Lens](#) in the *Amazon S3 User Guide*. For a complete list of S3 Storage Lens metrics, see [S3 Storage Lens metrics glossary](#) in the *Amazon S3 User Guide*.

Contents

BucketLevel

A container element for the S3 Storage Lens bucket-level configuration.

Type: [BucketLevel](#) data type

Required: Yes

ActivityMetrics

A container element for S3 Storage Lens activity metrics.

Type: [ActivityMetrics](#) data type

Required: No

AdvancedCostOptimizationMetrics

A container element for S3 Storage Lens advanced cost-optimization metrics.

Type: [AdvancedCostOptimizationMetrics](#) data type

Required: No

AdvancedDataProtectionMetrics

A container element for S3 Storage Lens advanced data-protection metrics.

Type: [AdvancedDataProtectionMetrics](#) data type

Required: No

AdvancedPerformanceMetrics

A container element for S3 Storage Lens advanced performance metrics.

Type: [AdvancedPerformanceMetrics](#) data type

Required: No

DetailedStatusCodesMetrics

A container element for detailed status code metrics.

Type: [DetailedStatusCodesMetrics](#) data type

Required: No

StorageLensGroupLevel

A container element for S3 Storage Lens groups metrics.

Type: [StorageLensGroupLevel](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ActivityMetrics

Service: Amazon S3 Control

The container element for Amazon S3 Storage Lens activity metrics. Activity metrics show details about how your storage is requested, such as requests (for example, All requests, Get requests, Put requests), bytes uploaded or downloaded, and errors.

For more information about S3 Storage Lens, see [Assessing your storage activity and usage with S3 Storage Lens](#) in the *Amazon S3 User Guide*. For a complete list of S3 Storage Lens metrics, see [S3 Storage Lens metrics glossary](#) in the *Amazon S3 User Guide*.

Contents

IsEnabled

A container that indicates whether activity metrics are enabled.

Type: Boolean

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

AdvancedCostOptimizationMetrics

Service: Amazon S3 Control

The container element for Amazon S3 Storage Lens advanced cost-optimization metrics. Advanced cost-optimization metrics provide insights that you can use to manage and optimize your storage costs, for example, lifecycle rule counts for transitions, expirations, and incomplete multipart uploads.

For more information about S3 Storage Lens, see [Assessing your storage activity and usage with S3 Storage Lens](#) in the *Amazon S3 User Guide*. For a complete list of S3 Storage Lens metrics, see [S3 Storage Lens metrics glossary](#) in the *Amazon S3 User Guide*.

Contents

IsEnabled

A container that indicates whether advanced cost-optimization metrics are enabled.

Type: Boolean

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

AdvancedDataProtectionMetrics

Service: Amazon S3 Control

The container element for Amazon S3 Storage Lens advanced data-protection metrics. Advanced data-protection metrics provide insights that you can use to perform audits and protect your data, for example replication rule counts within and across Regions.

For more information about S3 Storage Lens, see [Assessing your storage activity and usage with S3 Storage Lens](#) in the *Amazon S3 User Guide*. For a complete list of S3 Storage Lens metrics, see [S3 Storage Lens metrics glossary](#) in the *Amazon S3 User Guide*.

Contents

IsEnabled

A container that indicates whether advanced data-protection metrics are enabled.

Type: Boolean

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

AdvancedPerformanceMetrics

Service: Amazon S3 Control

The container element for S3 Storage Lens advanced performance metrics. Advanced performance metrics provide insights into application performance, such as request efficiency and access patterns. These metrics help you optimize your S3 storage for both cost and performance by providing detailed analytics on how your applications interact with S3 resources.

For more information about S3 Storage Lens, see [Assessing your storage activity and usage with S3 Storage Lens](#) in the *Amazon S3 User Guide*. For a complete list of S3 Storage Lens metrics, see [S3 Storage Lens metrics glossary](#) in the *Amazon S3 User Guide*.

Contents

IsEnabled

A container that indicates whether S3 Storage Lens advanced performance metrics are enabled.

Type: Boolean

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

AsyncErrorDetails

Service: Amazon S3 Control

Error details for the failed asynchronous operation.

Contents

Code

A string that uniquely identifies the error condition.

Type: String

Length Constraints: Maximum length of 1024.

Required: No

Message

A generic description of the error condition in English.

Type: String

Length Constraints: Maximum length of 1024.

Required: No

RequestId

The ID of the request associated with the error.

Type: String

Length Constraints: Maximum length of 1024.

Required: No

Resource

The identifier of the resource associated with the error.

Type: String

Length Constraints: Maximum length of 1024.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

AsyncOperation

Service: Amazon S3 Control

A container for the information about an asynchronous operation.

Contents

CreationTime

The time that the request was sent to the service.

Type: Timestamp

Required: No

Operation

The specific operation for the asynchronous request.

Type: String

Valid Values: CreateMultiRegionAccessPoint | DeleteMultiRegionAccessPoint | PutMultiRegionAccessPointPolicy

Required: No

RequestParameters

The parameters associated with the request.

Type: [AsyncRequestParameters](#) data type

Required: No

RequestStatus

The current status of the request.

Type: String

Required: No

RequestTokenARN

The request token associated with the request.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `arn: . +`

Required: No

ResponseDetails

The details of the response.

Type: [AsyncResponseDetails](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

AsyncRequestParameters

Service: Amazon S3 Control

A container for the request parameters associated with an asynchronous request.

Contents

CreateMultiRegionAccessPointRequest

A container of the parameters for a [CreateMultiRegionAccessPoint](#) request.

Type: [CreateMultiRegionAccessPointInput](#) data type

Required: No

DeleteMultiRegionAccessPointRequest

A container of the parameters for a [DeleteMultiRegionAccessPoint](#) request.

Type: [DeleteMultiRegionAccessPointInput](#) data type

Required: No

PutMultiRegionAccessPointPolicyRequest

A container of the parameters for a [PutMultiRegionAccessPoint](#) request.

Type: [PutMultiRegionAccessPointPolicyInput](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

AsyncResponseDetails

Service: Amazon S3 Control

A container for the response details that are returned when querying about an asynchronous request.

Contents

ErrorDetails

Error details for an asynchronous request.

Type: [AsyncErrorDetails](#) data type

Required: No

MultiRegionAccessPointDetails

The details for the Multi-Region Access Point.

Type: [MultiRegionAccessPointsAsyncResponse](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

AwsLambdaTransformation

Service: Amazon S3 Control

AWS Lambda function used to transform objects through an Object Lambda Access Point.

Contents

FunctionArn

The Amazon Resource Name (ARN) of the AWS Lambda function.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: (arn:(aws[a-zA-Z-]*)?:lambda:)?([a-z]{2}((-gov)|(-iso([a-z]?)))?-[a-z]+-\d{1}:)?(\d{12}:)?(function:)?([a-zA-Z0-9-_.+)(:(\$LATEST|[a-zA-Z0-9-_.+))?)?

Required: Yes

FunctionPayload

Additional JSON that provides supplemental data to the Lambda function used to transform objects.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

BucketLevel

Service: Amazon S3 Control

A container for the bucket-level configuration for Amazon S3 Storage Lens.

For more information about S3 Storage Lens, see [Assessing your storage activity and usage with S3 Storage Lens](#) in the *Amazon S3 User Guide*.

Contents

ActivityMetrics

A container for the bucket-level activity metrics for S3 Storage Lens.

Type: [ActivityMetrics](#) data type

Required: No

AdvancedCostOptimizationMetrics

A container for bucket-level advanced cost-optimization metrics for S3 Storage Lens.

Type: [AdvancedCostOptimizationMetrics](#) data type

Required: No

AdvancedDataProtectionMetrics

A container for bucket-level advanced data-protection metrics for S3 Storage Lens.

Type: [AdvancedDataProtectionMetrics](#) data type

Required: No

AdvancedPerformanceMetrics

A container for bucket-level advanced performance metrics for S3 Storage Lens.

Type: [AdvancedPerformanceMetrics](#) data type

Required: No

DetailedStatusCodesMetrics

A container for bucket-level detailed status code metrics for S3 Storage Lens.

Type: [DetailedStatusCodesMetrics](#) data type

Required: No

PrefixLevel

A container for the prefix-level metrics for S3 Storage Lens.

Type: [PrefixLevel](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

CloudWatchMetrics

Service: Amazon S3 Control

A container for enabling Amazon CloudWatch publishing for S3 Storage Lens metrics.

For more information about publishing S3 Storage Lens metrics to CloudWatch, see [Monitor S3 Storage Lens metrics in CloudWatch](#) in the *Amazon S3 User Guide*.

Contents

IsEnabled

A container that indicates whether CloudWatch publishing for S3 Storage Lens metrics is enabled. A value of `true` indicates that CloudWatch publishing for S3 Storage Lens metrics is enabled.

Type: Boolean

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

CreateBucketConfiguration

Service: Amazon S3 Control

The container for the bucket configuration.

Note

This is not supported by Amazon S3 on Outposts buckets.

Contents

LocationConstraint

Specifies the Region where the bucket will be created. If you are creating a bucket on the US East (N. Virginia) Region (us-east-1), you do not need to specify the location.

Note

This is not supported by Amazon S3 on Outposts buckets.

Type: String

Valid Values: EU | eu-west-1 | us-west-1 | us-west-2 | ap-south-1 | ap-southeast-1 | ap-southeast-2 | ap-northeast-1 | sa-east-1 | cn-north-1 | eu-central-1

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

CreateMultiRegionAccessPointInput

Service: Amazon S3 Control

A container for the information associated with a [CreateMultiRegionAccessPoint](#) request.

Contents

Name

The name of the Multi-Region Access Point associated with this request.

Type: String

Length Constraints: Maximum length of 50.

Pattern: `^[a-z0-9][-a-z0-9]{1,48}[a-z0-9]$`

Required: Yes

Regions

The buckets in different Regions that are associated with the Multi-Region Access Point.

Type: Array of [Region](#) data types

Required: Yes

PublicAccessBlock

The `PublicAccessBlock` configuration that you want to apply to this Amazon S3 account. You can enable the configuration options in any combination. For more information about when Amazon S3 considers a bucket or object public, see [The Meaning of "Public"](#) in the *Amazon S3 User Guide*.

This data type is not supported for Amazon S3 on Outposts.

Type: [PublicAccessBlockConfiguration](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Credentials

Service: Amazon S3 Control

The AWS Security Token Service temporary credential that S3 Access Grants vends to grantees and client applications.

Contents

AccessKeyId

The unique access key ID of the AWS STS temporary credential that S3 Access Grants vends to grantees and client applications.

Type: String

Required: No

Expiration

The expiration date and time of the temporary credential that S3 Access Grants vends to grantees and client applications.

Type: Timestamp

Required: No

SecretAccessKey

The secret access key of the AWS STS temporary credential that S3 Access Grants vends to grantees and client applications.

Type: String

Required: No

SessionToken

The AWS STS temporary credential that S3 Access Grants vends to grantees and client applications.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

DeleteMarkerReplication

Service: Amazon S3 Control

Specifies whether S3 on Outposts replicates delete markers. If you specify a `Filter` element in your replication configuration, you must also include a `DeleteMarkerReplication` element. If your `Filter` includes a `Tag` element, the `DeleteMarkerReplication` element's `Status` child element must be set to `Disabled`, because S3 on Outposts does not support replicating delete markers for tag-based rules.

For more information about delete marker replication, see [How delete operations affect replication](#) in the *Amazon S3 User Guide*.

Contents

Status

Indicates whether to replicate delete markers.

Type: String

Valid Values: `Enabled` | `Disabled`

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

DeleteMultiRegionAccessPointInput

Service: Amazon S3 Control

A container for the information associated with a [DeleteMultiRegionAccessPoint](#) request.

Contents

Name

The name of the Multi-Region Access Point associated with this request.

Type: String

Length Constraints: Maximum length of 50.

Pattern: `^[a-z0-9][-a-z0-9]{1,48}[a-z0-9]$`

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Destination

Service: Amazon S3 Control

Specifies information about the replication destination bucket and its settings for an S3 on Outposts replication configuration.

Contents

Bucket

The Amazon Resource Name (ARN) of the access point for the destination bucket where you want S3 on Outposts to store the replication results.

Type: String

Required: Yes

AccessControlTranslation

Specify this property only in a cross-account scenario (where the source and destination bucket owners are not the same), and you want to change replica ownership to the AWS account that owns the destination bucket. If this property is not specified in the replication configuration, the replicas are owned by same AWS account that owns the source object.

Note

This is not supported by Amazon S3 on Outposts buckets.

Type: [AccessControlTranslation](#) data type

Required: No

Account

The destination bucket owner's account ID.

Type: String

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: No

EncryptionConfiguration

A container that provides information about encryption. If `SourceSelectionCriteria` is specified, you must specify this element.

Note

This is not supported by Amazon S3 on Outposts buckets.

Type: [EncryptionConfiguration](#) data type

Required: No

Metrics

A container that specifies replication metrics-related settings.

Type: [Metrics](#) data type

Required: No

ReplicationTime

A container that specifies S3 Replication Time Control (S3 RTC) settings, including whether S3 RTC is enabled and the time when all objects and operations on objects must be replicated. Must be specified together with a `Metrics` block.

Note

This is not supported by Amazon S3 on Outposts buckets.

Type: [ReplicationTime](#) data type

Required: No

StorageClass

The storage class to use when replicating objects. All objects stored on S3 on Outposts are stored in the `OUTPOSTS` storage class. S3 on Outposts uses the `OUTPOSTS` storage class to create the object replicas.

Note

Values other than OUTPOSTS aren't supported by Amazon S3 on Outposts.

Type: String

Valid Values: STANDARD | REDUCED_REDUNDANCY | STANDARD_IA | ONEZONE_IA | INTELLIGENT_TIERING | GLACIER | DEEP_ARCHIVE | OUTPOSTS | GLACIER_IR

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

DetailedStatusCodesMetrics

Service: Amazon S3 Control

The container element for Amazon S3 Storage Lens detailed status code metrics. Detailed status code metrics generate metrics for HTTP status codes, such as 200 OK, 403 Forbidden, 503 Service Unavailable and others.

For more information about S3 Storage Lens, see [Assessing your storage activity and usage with S3 Storage Lens](#) in the *Amazon S3 User Guide*. For a complete list of S3 Storage Lens metrics, see [S3 Storage Lens metrics glossary](#) in the *Amazon S3 User Guide*.

Contents

IsEnabled

A container that indicates whether detailed status code metrics are enabled.

Type: Boolean

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

DSSEKMSFilter

Service: Amazon S3 Control

A filter that returns objects that are encrypted by dual-layer server-side encryption with AWS Key Management Service (KMS) keys (DSSE-KMS). You can further refine your filtering by optionally providing a KMS Key ARN to filter objects encrypted by a specific key.

Contents

KmsKeyArn

The Amazon Resource Name (ARN) of the customer managed KMS key to use for the filter to return objects that are encrypted by the specified key. For best performance, use keys in the same Region as the S3 Batch Operations job.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `arn:aws[a-zA-Z0-9-]*:kms:[a-z0-9-]+:[0-9]{12}:key/[a-zA-Z0-9-]+`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

EncryptionConfiguration

Service: Amazon S3 Control

Specifies encryption-related information for an Amazon S3 bucket that is a destination for replicated objects. If you're specifying a customer managed KMS key, we recommend using a fully qualified KMS key ARN. If you use a KMS key alias instead, then AWS KMS resolves the key within the requester's account. This behavior can result in data that's encrypted with a KMS key that belongs to the requester, and not the bucket owner.

Note

This is not supported by Amazon S3 on Outposts buckets.

Contents

ReplicaKmsKeyID

Specifies the ID of the customer managed AWS KMS key that's stored in AWS Key Management Service (AWS KMS) for the destination bucket. This ID is either the Amazon Resource Name (ARN) for the KMS key or the alias ARN for the KMS key. Amazon S3 uses this KMS key to encrypt replica objects. Amazon S3 supports only symmetric encryption KMS keys. For more information, see [Symmetric encryption KMS keys](#) in the *AWS Key Management Service Developer Guide*.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

EstablishedMultiRegionAccessPointPolicy

Service: Amazon S3 Control

The last established access control policy for a Multi-Region Access Point.

When you update the policy, the update is first listed as the proposed policy. After the update is finished and all Regions have been updated, the proposed policy is listed as the established policy. If both policies have the same version number, the proposed policy is the established policy.

Contents

Policy

The details of the last established policy.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Exclude

Service: Amazon S3 Control

A container for what Amazon S3 Storage Lens will exclude.

Contents

Buckets

A container for the S3 Storage Lens bucket excludes.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 128.

Pattern: `arn:[^:]+:s3:.*`

Required: No

Regions

A container for the S3 Storage Lens Region excludes.

Type: Array of strings

Length Constraints: Minimum length of 5. Maximum length of 30.

Pattern: `[a-z0-9\-]+`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ExistingObjectReplication

Service: Amazon S3 Control

An optional configuration to replicate existing source bucket objects.

Note

This is not supported by Amazon S3 on Outposts buckets.

Contents

Status

Specifies whether Amazon S3 replicates existing source bucket objects.

Type: String

Valid Values: Enabled | Disabled

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

GeneratedManifestEncryption

Service: Amazon S3 Control

The encryption configuration to use when storing the generated manifest.

Contents

SSEKMS

Configuration details on how SSE-KMS is used to encrypt generated manifest objects.

Type: [SSEKMSEncryption](#) data type

Required: No

SSES3

Specifies the use of SSE-S3 to encrypt generated manifest objects.

Type: [SSES3Encryption](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Grantee

Service: Amazon S3 Control

The user, group, or role to which you are granting access. You can grant access to an IAM user or role. If you have added your corporate directory to AWS IAM Identity Center and associated your Identity Center instance with your S3 Access Grants instance, the grantee can also be a corporate directory user or group.

Contents

GranteeIdentifier

The unique identifier of the Grantee. If the grantee type is IAM, the identifier is the IAM Amazon Resource Name (ARN) of the user or role. If the grantee type is a directory user or group, the identifier is 128-bit universally unique identifier (UUID) in the format a1b2c3d4-5678-90ab-cdef-EXAMPLE11111. You can obtain this UUID from your AWS IAM Identity Center instance.

Type: String

Required: No

GranteeType

The type of the grantee to which access has been granted. It can be one of the following values:

- IAM - An IAM user or role.
- DIRECTORY_USER - Your corporate directory user. You can use this option if you have added your corporate identity directory to IAM Identity Center and associated the IAM Identity Center instance with your S3 Access Grants instance.
- DIRECTORY_GROUP - Your corporate directory group. You can use this option if you have added your corporate identity directory to IAM Identity Center and associated the IAM Identity Center instance with your S3 Access Grants instance.

Type: String

Valid Values: DIRECTORY_USER | DIRECTORY_GROUP | IAM

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Include

Service: Amazon S3 Control

A container for what Amazon S3 Storage Lens configuration includes.

Contents

Buckets

A container for the S3 Storage Lens bucket includes.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 128.

Pattern: `arn:[^:]+:s3:.*`

Required: No

Regions

A container for the S3 Storage Lens Region includes.

Type: Array of strings

Length Constraints: Minimum length of 5. Maximum length of 30.

Pattern: `[a-z0-9\ - ]+`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

JobDescriptor

Service: Amazon S3 Control

A container element for the job configuration and status information returned by a `Describe Job` request.

Contents

ConfirmationRequired

Indicates whether confirmation is required before Amazon S3 begins running the specified job. Confirmation is required only for jobs created through the Amazon S3 console.

Type: Boolean

Required: No

CreationTime

A timestamp indicating when this job was created.

Type: Timestamp

Required: No

Description

The description for this job, if one was provided in this job's `Create Job` request.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Required: No

FailureReasons

If the specified job failed, this field contains information describing the failure.

Type: Array of [JobFailure](#) data types

Required: No

GeneratedManifestDescriptor

The attribute of the `JobDescriptor` containing details about the job's generated manifest.

Type: [S3GeneratedManifestDescriptor](#) data type

Required: No

JobArn

The Amazon Resource Name (ARN) for this job.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `arn:[^:]+:s3:[a-zA-Z0-9\-.]+\d{12}:job\/.*`

Required: No

JobId

The ID for the specified job.

Type: String

Length Constraints: Minimum length of 5. Maximum length of 36.

Pattern: `[a-zA-Z0-9\-\_]+`

Required: No

Manifest

The configuration information for the specified job's manifest object.

Type: [JobManifest](#) data type

Required: No

ManifestGenerator

The manifest generator that was used to generate a job manifest for this job.

Type: [JobManifestGenerator](#) data type

Note: This object is a Union. Only one member of this object can be specified or returned.

Required: No

Operation

The operation that the specified job is configured to run on the objects listed in the manifest.

Type: [JobOperation](#) data type

Required: No

Priority

The priority of the specified job.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 2147483647.

Required: No

ProgressSummary

Describes the total number of tasks that the specified job has run, the number of tasks that succeeded, and the number of tasks that failed.

Type: [JobProgressSummary](#) data type

Required: No

Report

Contains the configuration information for the job-completion report if you requested one in the `Create Job` request.

Type: [JobReport](#) data type

Required: No

RoleArn

The Amazon Resource Name (ARN) for the AWS Identity and Access Management (IAM) role assigned to run the tasks for this job.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:[^:]+:iam::\d{12}:role/.*`

Required: No

Status

The current status of the specified job.

Type: String

Valid Values: Active | Cancelled | Cancelling | Complete | Completing
| Failed | Failing | New | Paused | Pausing | Preparing | Ready |
Suspended

Required: No

StatusUpdateReason

The reason for updating the job.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Required: No

SuspendedCause

The reason why the specified job was suspended. A job is only suspended if you create it through the Amazon S3 console. When you create the job, it enters the Suspended state to await confirmation before running. After you confirm the job, it automatically exits the Suspended state.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

SuspendedDate

The timestamp when this job was suspended, if it has been suspended.

Type: Timestamp

Required: No

TerminationDate

A timestamp indicating when this job terminated. A job's termination date is the date and time when it succeeded, failed, or was canceled.

Type: Timestamp

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

JobFailure

Service: Amazon S3 Control

If this job failed, this element indicates why the job failed.

Contents

FailureCode

The failure code, if any, for the specified job.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Required: No

FailureReason

The failure reason, if any, for the specified job.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

JobListDescriptor

Service: Amazon S3 Control

Contains the configuration and status information for a single job retrieved as part of a job list.

Contents

CreationTime

A timestamp indicating when the specified job was created.

Type: Timestamp

Required: No

Description

The user-specified description that was included in the specified job's `Create Job` request.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Required: No

JobId

The ID for the specified job.

Type: String

Length Constraints: Minimum length of 5. Maximum length of 36.

Pattern: `[a-zA-Z0-9\-\_\+]`

Required: No

Operation

The operation that the specified job is configured to run on every object listed in the manifest.

Type: String

Valid Values: `LambdaInvoke` | `S3PutObjectCopy` | `S3PutObjectAcl` | `S3PutObjectTagging` | `S3DeleteObjectTagging` | `S3InitiateRestoreObject`

| S3PutObjectLegalHold | S3PutObjectRetention | S3ReplicateObject |
S3ComputeObjectChecksum | S3UpdateObjectEncryption

Required: No

Priority

The current priority for the specified job.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 2147483647.

Required: No

ProgressSummary

Describes the total number of tasks that the specified job has run, the number of tasks that succeeded, and the number of tasks that failed.

Type: [JobProgressSummary](#) data type

Required: No

Status

The specified job's current status.

Type: String

Valid Values: Active | Cancelled | Cancelling | Complete | Completing
| Failed | Failing | New | Paused | Pausing | Preparing | Ready |
Suspended

Required: No

TerminationDate

A timestamp indicating when the specified job terminated. A job's termination date is the date and time when it succeeded, failed, or was canceled.

Type: Timestamp

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

JobManifest

Service: Amazon S3 Control

Contains the configuration information for a job's manifest.

Contents

Location

Contains the information required to locate the specified job's manifest. Manifests can't be imported from directory buckets. For more information, see [Directory buckets](#).

Type: [JobManifestLocation](#) data type

Required: Yes

Spec

Describes the format of the specified job's manifest. If the manifest is in CSV format, also describes the columns contained within the manifest.

Type: [JobManifestSpec](#) data type

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

JobManifestGenerator

Service: Amazon S3 Control

Configures the type of the job's ManifestGenerator.

Contents

Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

S3JobManifestGenerator

The S3 job ManifestGenerator's configuration details.

Type: [S3JobManifestGenerator](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

JobManifestGeneratorFilter

Service: Amazon S3 Control

The filter used to describe a set of objects for the job's manifest.

Contents

CreatedAfter

If provided, the generated manifest includes only source bucket objects that were created after this time.

Type: Timestamp

Required: No

CreatedBefore

If provided, the generated manifest includes only source bucket objects that were created before this time.

Type: Timestamp

Required: No

EligibleForReplication

Include objects in the generated manifest only if they are eligible for replication according to the Replication configuration on the source bucket.

Type: Boolean

Required: No

KeyNameConstraint

If provided, the generated manifest includes only source bucket objects whose object keys match the string constraints specified for `MatchAnyPrefix`, `MatchAnySuffix`, and `MatchAnySubstring`.

Type: [KeyNameConstraint](#) data type

Required: No

MatchAnyObjectEncryption

If provided, the generated object list includes only source bucket objects with the indicated server-side encryption type (SSE-S3, SSE-KMS, DSSE-KMS, SSE-C, or NOT-SSE).

Type: Array of [ObjectEncryptionFilter](#) data types

Array Members: Fixed number of 1 item.

Required: No

MatchAnyStorageClass

If provided, the generated manifest includes only source bucket objects that are stored with the specified storage class.

Type: Array of strings

Valid Values: STANDARD | STANDARD_IA | ONEZONE_IA | GLACIER |
INTELLIGENT_TIERING | DEEP_ARCHIVE | GLACIER_IR

Required: No

ObjectReplicationStatuses

If provided, the generated manifest includes only source bucket objects that have one of the specified Replication statuses.

Type: Array of strings

Valid Values: COMPLETED | FAILED | REPLICA | NONE

Required: No

ObjectSizeGreaterThanBytes

If provided, the generated manifest includes only source bucket objects whose file size is greater than the specified number of bytes.

Type: Long

Required: No

ObjectSizeLessThanBytes

If provided, the generated manifest includes only source bucket objects whose file size is less than the specified number of bytes.

Type: Long

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

JobManifestLocation

Service: Amazon S3 Control

Contains the information required to locate a manifest object. Manifests can't be imported from directory buckets. For more information, see [Directory buckets](#).

Contents

ETag

The ETag for the specified manifest object.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

ObjectArn

The Amazon Resource Name (ARN) for a manifest object.

Important

When you're using XML requests, you must replace special characters (such as carriage returns) in object keys with their equivalent XML entity codes. For more information, see [XML-related object key constraints](#) in the *Amazon S3 User Guide*.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2000.

Pattern: `arn:[^:]+:s3:.*`

Required: Yes

ObjectVersionId

The optional version ID to identify a specific version of the manifest object.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2000.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

JobManifestSpec

Service: Amazon S3 Control

Describes the format of a manifest. If the manifest is in CSV format, also describes the columns contained within the manifest.

Contents

Format

Indicates which of the available formats the specified manifest uses.

Type: String

Valid Values: S3BatchOperations_CSV_20180820 |
S3InventoryReport_CSV_20161130

Required: Yes

Fields

If the specified manifest object is in the S3BatchOperations_CSV_20180820 format, this element describes which columns contain the required data.

Type: Array of strings

Valid Values: Ignore | Bucket | Key | VersionId

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

JobOperation

Service: Amazon S3 Control

The operation that you want this job to perform on every object listed in the manifest. For more information about the available operations, see [Operations](#) in the *Amazon S3 User Guide*.

Contents

LambdaInvoke

Directs the specified job to invoke an AWS Lambda function on every object in the manifest.

Type: [LambdaInvokeOperation](#) data type

Required: No

S3ComputeObjectChecksum

Directs the specified job to compute checksum values for every object in the manifest.

Type: [S3ComputeObjectChecksumOperation](#) data type

Required: No

S3DeleteObjectTagging

Directs the specified job to execute a DELETE Object tagging call on every object in the manifest.

Note

This functionality is not supported by directory buckets.

Type: [S3DeleteObjectTaggingOperation](#) data type

Required: No

S3InitiateRestoreObject

Directs the specified job to initiate restore requests for every archived object in the manifest.

Note

This functionality is not supported by directory buckets.

Type: [S3InitiateRestoreObjectOperation](#) data type

Required: No

S3PutObjectAcl

Directs the specified job to run a PutObjectAcl call on every object in the manifest.

Note

This functionality is not supported by directory buckets.

Type: [S3SetObjectAclOperation](#) data type

Required: No

S3PutObjectCopy

Directs the specified job to run a PUT Copy object call on every object in the manifest.

Type: [S3CopyObjectOperation](#) data type

Required: No

S3PutObjectLegalHold

Contains the configuration for an S3 Object Lock legal hold operation that an S3 Batch Operations job passes to every object to the underlying PutObjectLegalHold API operation. For more information, see [Using S3 Object Lock legal hold with S3 Batch Operations](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported by directory buckets.

Type: [S3SetObjectLegalHoldOperation](#) data type

Required: No

S3PutObjectRetention

Contains the configuration parameters for the Object Lock retention action for an S3 Batch Operations job. Batch Operations passes every object to the underlying PutObjectRetention API operation. For more information, see [Using S3 Object Lock retention with S3 Batch Operations](#) in the *Amazon S3 User Guide*.

 Note

This functionality is not supported by directory buckets.

Type: [S3SetObjectRetentionOperation](#) data type

Required: No

S3PutObjectTagging

Directs the specified job to run a PUT Object tagging call on every object in the manifest.

 Note

This functionality is not supported by directory buckets.

Type: [S3SetObjectTaggingOperation](#) data type

Required: No

S3ReplicateObject

Directs the specified job to invoke ReplicateObject on every object in the job's manifest.

 Note

This functionality is not supported by directory buckets.

Type: [S3ReplicateObjectOperation](#) data type

Required: No

S3UpdateObjectEncryption

Updates the server-side encryption type of an existing encrypted object in a general purpose bucket. You can use the `UpdateObjectEncryption` operation to change encrypted objects from server-side encryption with Amazon S3 managed keys (SSE-S3) to server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS), or to apply S3 Bucket Keys. You can also use the `UpdateObjectEncryption` operation to change the customer-managed KMS key used to encrypt your data so that you can comply with custom key-rotation standards.

Type: [S3UpdateObjectEncryptionOperation](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

JobProgressSummary

Service: Amazon S3 Control

Describes the total number of tasks that the specified job has started, the number of tasks that succeeded, and the number of tasks that failed.

Contents

NumberOfTasksFailed

Type: Long

Valid Range: Minimum value of 0.

Required: No

NumberOfTasksSucceeded

Type: Long

Valid Range: Minimum value of 0.

Required: No

Timers

The JobTimers attribute of a job's progress summary.

Type: [JobTimers](#) data type

Required: No

TotalNumberOfTasks

Type: Long

Valid Range: Minimum value of 0.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

JobReport

Service: Amazon S3 Control

Contains the configuration parameters for a job-completion report.

Contents

Enabled

Indicates whether the specified job will generate a job-completion report.

Type: Boolean

Required: Yes

Bucket

The Amazon Resource Name (ARN) for the bucket where specified job-completion report will be stored.

Note

Directory buckets - Directory buckets aren't supported as a location for Batch Operations to store job completion reports.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 128.

Pattern: `arn:[^:]+:s3:.*`

Required: No

ExpectedBucketOwner

Lists the AWS account ID that owns the target bucket, where the completion report is received.

Type: String

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: No

Format

The format of the specified job-completion report.

Type: String

Valid Values: Report_CSV_20180820

Required: No

Prefix

An optional prefix to describe where in the specified bucket the job-completion report will be stored. Amazon S3 stores the job-completion report at <prefix>/job-<job-id>/report.json.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Required: No

ReportScope

Indicates whether the job-completion report will include details of all tasks or only failed tasks.

Type: String

Valid Values: AllTasks | FailedTasksOnly

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

JobTimers

Service: Amazon S3 Control

Provides timing details for the job.

Contents

ElapsedTimeInActiveSeconds

Indicates the elapsed time in seconds the job has been in the Active job state.

Type: Long

Valid Range: Minimum value of 0.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

KeyNameConstraint

Service: Amazon S3 Control

If provided, the generated manifest includes only source bucket objects whose object keys match the string constraints specified for `MatchAnyPrefix`, `MatchAnySuffix`, and `MatchAnySubstring`.

Contents

MatchAnyPrefix

If provided, the generated manifest includes objects where the specified string appears at the start of the object key string. Each `KeyNameConstraint` filter accepts an array of strings with a length of 1 string.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

MatchAnySubstring

If provided, the generated manifest includes objects where the specified string appears anywhere within the object key string. Each `KeyNameConstraint` filter accepts an array of strings with a length of 1 string.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

MatchAnySuffix

If provided, the generated manifest includes objects where the specified string appears at the end of the object key string. Each `KeyNameConstraint` filter accepts an array of strings with a length of 1 string.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

LambdaInvokeOperation

Service: Amazon S3 Control

Contains the configuration parameters for a Lambda Invoke operation.

Contents

FunctionArn

The Amazon Resource Name (ARN) for the AWS Lambda function that the specified job will invoke on every object in the manifest.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: (arn:(aws[a-zA-Z-]*)?:lambda:)?([a-z]{2}((-gov)|(-iso([a-z]?)))?-[a-z]+-\d{1}:)?(\d{12}:)?(function:)?([a-zA-Z0-9-_.+)(:(\$LATEST|[a-zA-Z0-9-_.+))?)?

Required: No

InvocationSchemaVersion

Specifies the schema version for the payload that Batch Operations sends when invoking an AWS Lambda function. Version 1.0 is the default. Version 2.0 is required when you use Batch Operations to invoke AWS Lambda functions that act on directory buckets, or if you need to specify `UserArguments`. For more information, see [Automate object processing in Amazon S3 directory buckets with S3 Batch Operations and AWS Lambda](#) in the *AWS Storage Blog*.

Important

Ensure that your AWS Lambda function code expects `InvocationSchemaVersion 2.0` and uses bucket name rather than bucket ARN. If the `InvocationSchemaVersion` does not match what your AWS Lambda function expects, your function might not work as expected.

Note

Directory buckets - To initiate AWS Lambda function to perform custom actions on objects in directory buckets, you must specify 2.0.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Required: No

UserArguments

Key-value pairs that are passed in the payload that Batch Operations sends when invoking an AWS Lambda function. You must specify `InvocationSchemaVersion 2.0` for `LambdaInvoke` operations that include `UserArguments`. For more information, see [Automate object processing in Amazon S3 directory buckets with S3 Batch Operations and AWS Lambda](#) in the *AWS Storage Blog*.

Type: String to string map

Map Entries: Maximum number of 10 items.

Key Length Constraints: Minimum length of 1. Maximum length of 64.

Value Length Constraints: Maximum length of 1024.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

LifecycleConfiguration

Service: Amazon S3 Control

The container for the Outposts bucket lifecycle configuration.

Contents

Rules

A lifecycle rule for individual objects in an Outposts bucket.

Type: Array of [LifecycleRule](#) data types

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

LifecycleExpiration

Service: Amazon S3 Control

The container of the Outposts bucket lifecycle expiration.

Contents

Date

Indicates at what date the object is to be deleted. Should be in GMT ISO 8601 format.

Type: Timestamp

Required: No

Days

Indicates the lifetime, in days, of the objects that are subject to the rule. The value must be a non-zero positive integer.

Type: Integer

Required: No

ExpiredObjectDeleteMarker

Indicates whether Amazon S3 will remove a delete marker with no noncurrent versions. If set to true, the delete marker will be expired. If set to false, the policy takes no action. This cannot be specified with Days or Date in a Lifecycle Expiration Policy. To learn more about delete markers, see [Working with delete markers](#).

Type: Boolean

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)

- [AWS SDK for Ruby V3](#)

LifecycleRule

Service: Amazon S3 Control

The container for the Outposts bucket lifecycle rule.

Contents

Status

If 'Enabled', the rule is currently being applied. If 'Disabled', the rule is not currently being applied.

Type: String

Valid Values: Enabled | Disabled

Required: Yes

AbortIncompleteMultipartUpload

Specifies the days since the initiation of an incomplete multipart upload that Amazon S3 waits before permanently removing all parts of the upload. For more information, see [Aborting Incomplete Multipart Uploads Using a Bucket Lifecycle Configuration](#) in the *Amazon S3 User Guide*.

Type: [AbortIncompleteMultipartUpload](#) data type

Required: No

Expiration

Specifies the expiration for the lifecycle of the object in the form of date, days and, whether the object has a delete marker.

Type: [LifecycleExpiration](#) data type

Required: No

Filter

The container for the filter of lifecycle rule.

Type: [LifecycleRuleFilter](#) data type

Required: No

ID

Unique identifier for the rule. The value cannot be longer than 255 characters.

Type: String

Required: No

NoncurrentVersionExpiration

The noncurrent version expiration of the lifecycle rule.

Type: [NoncurrentVersionExpiration](#) data type

Required: No

NoncurrentVersionTransitions

Specifies the transition rule for the lifecycle rule that describes when noncurrent objects transition to a specific storage class. If your bucket is versioning-enabled (or versioning is suspended), you can set this action to request that Amazon S3 transition noncurrent object versions to a specific storage class at a set period in the object's lifetime.

Note

This is not supported by Amazon S3 on Outposts buckets.

Type: Array of [NoncurrentVersionTransition](#) data types

Required: No

Transitions

Specifies when an Amazon S3 object transitions to a specified storage class.

Note

This is not supported by Amazon S3 on Outposts buckets.

Type: Array of [Transition](#) data types

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

LifecycleRuleAndOperator

Service: Amazon S3 Control

The container for the Outposts bucket lifecycle rule and operator.

Contents

ObjectSizeGreaterThan

The non-inclusive minimum object size for the lifecycle rule. Setting this property to 7 means the rule applies to objects with a size that is greater than 7.

Type: Long

Required: No

ObjectSizeLessThan

The non-inclusive maximum object size for the lifecycle rule. Setting this property to 77 means the rule applies to objects with a size that is less than 77.

Type: Long

Required: No

Prefix

Prefix identifying one or more objects to which the rule applies.

Type: String

Required: No

Tags

All of these tags must exist in the object's tag set in order for the rule to apply.

Type: Array of [S3Tag](#) data types

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

LifecycleRuleFilter

Service: Amazon S3 Control

The container for the filter of the lifecycle rule.

Contents

And

The container for the AND condition for the lifecycle rule.

Type: [LifecycleRuleAndOperator](#) data type

Required: No

ObjectSizeGreaterThan

Minimum object size to which the rule applies.

Type: Long

Required: No

ObjectSizeLessThan

Maximum object size to which the rule applies.

Type: Long

Required: No

Prefix

Prefix identifying one or more objects to which the rule applies.

Important

When you're using XML requests, you must replace special characters (such as carriage returns) in object keys with their equivalent XML entity codes. For more information, see [XML-related object key constraints](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

Tag

A container for a key-value name pair.

Type: [S3Tag](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ListAccessGrantEntry

Service: Amazon S3 Control

Information about the access grant.

Contents

AccessGrantArn

The Amazon Resource Name (ARN) of the access grant.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:[a-z\-\+]:s3:[a-z0-9\-\+]:\d{12}:access\-\grants\grant/[a-zA-Z0-9\-\+]`

Required: No

AccessGrantId

The ID of the access grant. S3 Access Grants auto-generates this ID when you create the access grant.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\+]`

Required: No

AccessGrantsLocationConfiguration

The configuration options of the grant location. The grant location is the S3 path to the data to which you are granting access.

Type: [AccessGrantsLocationConfiguration](#) data type

Required: No

AccessGrantsLocationId

The ID of the registered location to which you are granting access. S3 Access Grants assigns this ID when you register the location. S3 Access Grants assigns the ID default to the default location `s3://` and assigns an auto-generated ID to other locations that you register.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\ - ]+`

Required: No

ApplicationArn

The Amazon Resource Name (ARN) of an AWS IAM Identity Center application associated with your Identity Center instance. If the grant includes an application ARN, the grantee can only access the S3 data through this application.

Type: String

Length Constraints: Minimum length of 10. Maximum length of 1224.

Pattern: `arn:[^:]+:sso:\d{12}:application/.*$`

Required: No

CreatedAt

The date and time when you created the S3 Access Grants instance.

Type: Timestamp

Required: No

Grantee

The user, group, or role to which you are granting access. You can grant access to an IAM user or role. If you have added your corporate directory to AWS IAM Identity Center and associated your Identity Center instance with your S3 Access Grants instance, the grantee can also be a corporate directory user or group.

Type: [Grantee](#) data type

Required: No

GrantScope

The S3 path of the data to which you are granting access. It is the result of appending the `Subprefix` to the location scope.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2000.

Pattern: `^\.+$`

Required: No

Permission

The type of access granted to your S3 data, which can be set to one of the following values:

- `READ` – Grant read-only access to the S3 data.
- `WRITE` – Grant write-only access to the S3 data.
- `READWRITE` – Grant both read and write access to the S3 data.

Type: String

Valid Values: `READ` | `WRITE` | `READWRITE`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ListAccessGrantsInstanceEntry

Service: Amazon S3 Control

Information about the S3 Access Grants instance.

Contents

AccessGrantsInstanceArn

The Amazon Resource Name (ARN) of the S3 Access Grants instance.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:[a-z\-\-]+:s3:[a-z0-9\-\-]+:\d{12}:access\-\grants\[a-zA-Z0-9\-\-]+`

Required: No

AccessGrantsInstanceId

The ID of the S3 Access Grants instance. The ID is default. You can have one S3 Access Grants instance per Region per account.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\-]+`

Required: No

CreatedAt

The date and time when you created the S3 Access Grants instance.

Type: Timestamp

Required: No

IdentityCenterApplicationArn

If you associated your S3 Access Grants instance with an AWS IAM Identity Center instance, this field returns the Amazon Resource Name (ARN) of the IAM Identity Center instance application;

a subresource of the original Identity Center instance. S3 Access Grants creates this Identity Center application for the specific S3 Access Grants instance.

Type: String

Length Constraints: Minimum length of 10. Maximum length of 1224.

Pattern: `arn:[^:]+:sso::\d{12}:application/.*$`

Required: No

IdentityCenterArn

This member has been deprecated.

If you associated your S3 Access Grants instance with an AWS IAM Identity Center instance, this field returns the Amazon Resource Name (ARN) of the IAM Identity Center instance application; a subresource of the original Identity Center instance. S3 Access Grants creates this Identity Center application for the specific S3 Access Grants instance.

Type: String

Length Constraints: Minimum length of 10. Maximum length of 1224.

Pattern: `arn:[^:]+:sso::(\d{12}){0,1}:instance/.*$`

Required: No

IdentityCenterInstanceArn

The Amazon Resource Name (ARN) of the AWS IAM Identity Center instance that you are associating with your S3 Access Grants instance. An IAM Identity Center instance is your corporate identity directory that you added to the IAM Identity Center. You can use the [ListInstances](#) API operation to retrieve a list of your Identity Center instances and their ARNs.

Type: String

Length Constraints: Minimum length of 10. Maximum length of 1224.

Pattern: `arn:[^:]+:sso::(\d{12}){0,1}:instance/.*$`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ListAccessGrantsLocationsEntry

Service: Amazon S3 Control

A container for information about the registered location.

Contents

AccessGrantsLocationArn

The Amazon Resource Name (ARN) of the registered location.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:[a-z\-\+]:s3:[a-z0-9\-\+]:\d{12}:access\-\grants\location/[a-zA-Z0-9\-\+]`

Required: No

AccessGrantsLocationId

The ID of the registered location to which you are granting access. S3 Access Grants assigns this ID when you register the location. S3 Access Grants assigns the ID default to the default location `s3://` and assigns an auto-generated ID to other locations that you register.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\+]`

Required: No

CreatedAt

The date and time when you registered the location.

Type: Timestamp

Required: No

IAMRoleArn

The Amazon Resource Name (ARN) of the IAM role for the registered location. S3 Access Grants assumes this role to manage access to the registered location.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:[^:]+:iam::\d{12}:role/.*`

Required: No

LocationScope

The S3 path to the location that you are registering. The location scope can be the default S3 location `s3://`, the S3 path to a bucket `s3://<bucket>`, or the S3 path to a bucket and prefix `s3://<bucket>/<prefix>`. A prefix in S3 is a string of characters at the beginning of an object key name used to organize the objects that you store in your S3 buckets. For example, object key names that start with the `engineering/` prefix or object key names that start with the `marketing/campaigns/` prefix.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2000.

Pattern: `^\.+$.+`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ListCallerAccessGrantsEntry

Service: Amazon S3 Control

Part of `ListCallerAccessGrantsResult`. Each entry includes the permission level (READ, WRITE, or READWRITE) and the grant scope of the access grant. If the grant also includes an application ARN, the grantee can only access the S3 data through this application.

Contents

ApplicationArn

The Amazon Resource Name (ARN) of an AWS IAM Identity Center application associated with your Identity Center instance. If the grant includes an application ARN, the grantee can only access the S3 data through this application.

Type: String

Length Constraints: Minimum length of 10. Maximum length of 1224.

Pattern: `arn:[^:]+:sso:\d{12}:application/.*$`

Required: No

GrantScope

The S3 path of the data to which you have been granted access.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2000.

Pattern: `^\.+$`

Required: No

Permission

The type of permission granted, which can be one of the following values:

- READ - Grants read-only access to the S3 data.
- WRITE - Grants write-only access to the S3 data.
- READWRITE - Grants both read and write access to the S3 data.

Type: String

Valid Values: READ | WRITE | READWRITE

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ListStorageLensConfigurationEntry

Service: Amazon S3 Control

Part of `ListStorageLensConfigurationResult`. Each entry includes the description of the S3 Storage Lens configuration, its home Region, whether it is enabled, its Amazon Resource Name (ARN), and config ID.

Contents

HomeRegion

A container for the S3 Storage Lens home Region. Your metrics data is stored and retained in your designated S3 Storage Lens home Region.

Type: String

Length Constraints: Minimum length of 5. Maximum length of 30.

Pattern: `[a-z0-9\-\-]+`

Required: Yes

Id

A container for the S3 Storage Lens configuration ID.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\_\.\-]+`

Required: Yes

StorageLensArn

The ARN of the S3 Storage Lens configuration. This property is read-only.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `arn:[a-z\-\-]+:s3:[a-z0-9\-\-]+:\d{12}:storage\-\-lens\-.*`

Required: Yes

IsEnabled

A container for whether the S3 Storage Lens configuration is enabled. This property is required.

Type: Boolean

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ListStorageLensGroupEntry

Service: Amazon S3 Control

Each entry contains a Storage Lens group that exists in the specified home Region.

Contents

HomeRegion

Contains the AWS Region where the Storage Lens group was created.

Type: String

Length Constraints: Minimum length of 5. Maximum length of 30.

Pattern: `[a-z0-9\-\_]+`

Required: Yes

Name

Contains the name of the Storage Lens group that exists in the specified home Region.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\_\_]+`

Required: Yes

StorageLensGroupArn

Contains the Amazon Resource Name (ARN) of the Storage Lens group. This property is read-only.

Type: String

Length Constraints: Minimum length of 4. Maximum length of 1024.

Pattern: `arn:[a-z\-\_]+:s3:[a-z0-9\-\_]+:\d{12}:storage\-\lens\-\group\/*`

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

MatchObjectAge

Service: Amazon S3 Control

A filter condition that specifies the object age range of included objects in days. Only integers are supported.

Contents

DaysGreaterThan

Specifies the maximum object age in days. Must be a positive whole number, greater than the minimum object age and less than or equal to 2,147,483,647.

Type: Integer

Required: No

DaysLessThan

Specifies the minimum object age in days. The value must be a positive whole number, greater than 0 and less than or equal to 2,147,483,647.

Type: Integer

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

MatchObjectSize

Service: Amazon S3 Control

A filter condition that specifies the object size range of included objects in bytes. Only integers are supported.

Contents

BytesGreaterThan

Specifies the minimum object size in Bytes. The value must be a positive number, greater than 0 and less than 50 TB.

Type: Long

Required: No

BytesLessThan

Specifies the maximum object size in Bytes. The value must be a positive number, greater than the minimum object size and less than 50 TB.

Type: Long

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Metrics

Service: Amazon S3 Control

A container that specifies replication metrics-related settings.

Contents

Status

Specifies whether replication metrics are enabled.

Type: String

Valid Values: Enabled | Disabled

Required: Yes

EventThreshold

A container that specifies the time threshold for emitting the `s3:Replication:OperationMissedThreshold` event.

Note

This is not supported by Amazon S3 on Outposts buckets.

Type: [ReplicationTimeValue](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

MultiRegionAccessPointPolicyDocument

Service: Amazon S3 Control

The Multi-Region Access Point access control policy.

When you update the policy, the update is first listed as the proposed policy. After the update is finished and all Regions have been updated, the proposed policy is listed as the established policy. If both policies have the same version number, the proposed policy is the established policy.

Contents

Established

The last established policy for the Multi-Region Access Point.

Type: [EstablishedMultiRegionAccessPointPolicy](#) data type

Required: No

Proposed

The proposed policy for the Multi-Region Access Point.

Type: [ProposedMultiRegionAccessPointPolicy](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

MultiRegionAccessPointRegionalResponse

Service: Amazon S3 Control

Status information for a single Multi-Region Access Point Region.

Contents

Name

The name of the Region in the Multi-Region Access Point.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Required: No

RequestStatus

The current status of the Multi-Region Access Point in this Region.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

MultiRegionAccessPointReport

Service: Amazon S3 Control

A collection of statuses for a Multi-Region Access Point in the various Regions it supports.

Contents

Alias

The alias for the Multi-Region Access Point. For more information about the distinction between the name and the alias of an Multi-Region Access Point, see [Rules for naming Amazon S3 Multi-Region Access Points](#).

Type: String

Length Constraints: Maximum length of 63.

Pattern: `^[a-z][a-z0-9]*[.]mrp$`

Required: No

CreatedAt

When the Multi-Region Access Point create request was received.

Type: Timestamp

Required: No

Name

The name of the Multi-Region Access Point.

Type: String

Length Constraints: Maximum length of 50.

Pattern: `^[a-z0-9] [-a-z0-9]{1,48} [a-z0-9]$`

Required: No

PublicAccessBlock

The `PublicAccessBlock` configuration that you want to apply to this Amazon S3 account. You can enable the configuration options in any combination. For more information about when

Amazon S3 considers a bucket or object public, see [The Meaning of "Public"](#) in the *Amazon S3 User Guide*.

This data type is not supported for Amazon S3 on Outposts.

Type: [PublicAccessBlockConfiguration](#) data type

Required: No

Regions

A collection of the Regions and buckets associated with the Multi-Region Access Point.

Type: Array of [RegionReport](#) data types

Required: No

Status

The current status of the Multi-Region Access Point.

CREATING and DELETING are temporary states that exist while the request is propagating and being completed. If a Multi-Region Access Point has a status of PARTIALLY_CREATED, you can retry creation or send a request to delete the Multi-Region Access Point. If a Multi-Region Access Point has a status of PARTIALLY_DELETED, you can retry a delete request to finish the deletion of the Multi-Region Access Point.

Type: String

Valid Values: READY | INCONSISTENT_ACROSS_REGIONS | CREATING | PARTIALLY_CREATED | PARTIALLY_DELETED | DELETING

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

MultiRegionAccessPointRoute

Service: Amazon S3 Control

A structure for a Multi-Region Access Point that indicates where Amazon S3 traffic can be routed. Routes can be either active or passive. Active routes can process Amazon S3 requests through the Multi-Region Access Point, but passive routes are not eligible to process Amazon S3 requests.

Each route contains the Amazon S3 bucket name and the AWS Region that the bucket is located in. The route also includes the `TrafficDialPercentage` value, which shows whether the bucket and Region are active (indicated by a value of 100) or passive (indicated by a value of 0).

Contents

TrafficDialPercentage

The traffic state for the specified bucket or AWS Region.

A value of 0 indicates a passive state, which means that no new traffic will be routed to the Region.

A value of 100 indicates an active state, which means that traffic will be routed to the specified Region.

When the routing configuration for a Region is changed from active to passive, any in-progress operations (uploads, copies, deletes, and so on) to the formerly active Region will continue to run until a final success or failure status is reached.

If all Regions in the routing configuration are designated as passive, you'll receive an `InvalidRequest` error.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 100.

Required: Yes

Bucket

The name of the Amazon S3 bucket for which you'll submit a routing configuration change. Either the `Bucket` or the `Region` value must be provided. If both are provided, the bucket must be in the specified Region.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: No

Region

The AWS Region to which you'll be submitting a routing configuration change. Either the Bucket or the Region value must be provided. If both are provided, the bucket must be in the specified Region.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

MultiRegionAccessPointsAsyncResponse

Service: Amazon S3 Control

The Multi-Region Access Point details that are returned when querying about an asynchronous request.

Contents

Regions

A collection of status information for the different Regions that a Multi-Region Access Point supports.

Type: Array of [MultiRegionAccessPointRegionalResponse](#) data types

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

NoncurrentVersionExpiration

Service: Amazon S3 Control

The container of the noncurrent version expiration.

Contents

NewerNoncurrentVersions

Specifies how many noncurrent versions S3 on Outposts will retain. If there are this many more recent noncurrent versions, S3 on Outposts will take the associated action. For more information about noncurrent versions, see [Lifecycle configuration elements](#) in the *Amazon S3 User Guide*.

Type: Integer

Required: No

NoncurrentDays

Specifies the number of days an object is noncurrent before Amazon S3 can perform the associated action. For information about the noncurrent days calculations, see [How Amazon S3 Calculates When an Object Became Noncurrent](#) in the *Amazon S3 User Guide*.

Type: Integer

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

NoncurrentVersionTransition

Service: Amazon S3 Control

The container for the noncurrent version transition.

Contents

NoncurrentDays

Specifies the number of days an object is noncurrent before Amazon S3 can perform the associated action. For information about the noncurrent days calculations, see [How Amazon S3 Calculates How Long an Object Has Been Noncurrent](#) in the *Amazon S3 User Guide*.

Type: Integer

Required: No

StorageClass

The class of storage used to store the object.

Type: String

Valid Values: GLACIER | STANDARD_IA | ONEZONE_IA | INTELLIGENT_TIERING | DEEP_ARCHIVE

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

NotSSEFilter

Service: Amazon S3 Control

A filter that returns objects that aren't server-side encrypted.

Contents

The members of this exception structure are context-dependent.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ObjectEncryption

Service: Amazon S3 Control

The updated server-side encryption type for this object. The `UpdateObjectEncryption` operation supports the SSE-KMS encryption type.

Valid Values: SSEKMS

Contents

SSEKMS

Specifies to update the object encryption type to server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS).

Type: [S3UpdateObjectEncryptionSSEKMS](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ObjectEncryptionFilter

Service: Amazon S3 Control

An optional filter for the S3JobManifestGenerator that identifies the subset of objects by encryption type.

Contents

Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

DSSEKMS

Filters for objects that are encrypted by dual-layer server-side encryption with AWS Key Management Service (KMS) keys (DSSE-KMS).

Type: [DSSEKMSFilter](#) data type

Required: No

NOTSSE

Filters for objects that are not encrypted by server-side encryption.

Type: [NotSSEFilter](#) data type

Required: No

SSEC

Filters for objects that are encrypted by server-side encryption with customer-provided keys (SSE-C).

Type: [SSECFilter](#) data type

Required: No

SSEKMS

Filters for objects that are encrypted by server-side encryption with AWS Key Management Service (KMS) keys (SSE-KMS).

Type: [SSEKMSFilter](#) data type

Required: No

SSES3

Filters for objects that are encrypted by server-side encryption with Amazon S3 managed keys (SSE-S3).

Type: [SSES3Filter](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ObjectLambdaAccessPoint

Service: Amazon S3 Control

An access point with an attached AWS Lambda function used to access transformed data from an Amazon S3 bucket.

Contents

Name

The name of the Object Lambda Access Point.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 45.

Pattern: `^[a-z0-9]([a-z0-9\-*][a-z0-9])?$`

Required: Yes

Alias

The alias of the Object Lambda Access Point.

Type: [ObjectLambdaAccessPointAlias](#) data type

Required: No

ObjectLambdaAccessPointArn

Specifies the ARN for the Object Lambda Access Point.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:[^:]+:s3-object-lambda:[^:]*:\d{12}:accesspoint/.*`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ObjectLambdaAccessPointAlias

Service: Amazon S3 Control

The alias of an Object Lambda Access Point. For more information, see [How to use a bucket-style alias for your S3 bucket Object Lambda Access Point](#).

Contents

Status

The status of the Object Lambda Access Point alias. If the status is `PROVISIONING`, the Object Lambda Access Point is provisioning the alias and the alias is not ready for use yet. If the status is `READY`, the Object Lambda Access Point alias is successfully provisioned and ready for use.

Type: String

Length Constraints: Minimum length of 2. Maximum length of 16.

Valid Values: `PROVISIONING` | `READY`

Required: No

Value

The alias value of the Object Lambda Access Point.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Pattern: `^[0-9a-z\\-]{3,63}`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)

- [AWS SDK for Ruby V3](#)

ObjectLambdaConfiguration

Service: Amazon S3 Control

A configuration used when creating an Object Lambda Access Point.

Contents

SupportingAccessPoint

Standard access point associated with the Object Lambda Access Point.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `arn:[^:]+:s3:[^:]*:\d{12}:accesspoint/.*`

Required: Yes

TransformationConfigurations

A container for transformation configurations for an Object Lambda Access Point.

Type: Array of [ObjectLambdaTransformationConfiguration](#) data types

Required: Yes

AllowedFeatures

A container for allowed features. Valid inputs are `GetObject-Range`, `GetObject-PartNumber`, `HeadObject-Range`, and `HeadObject-PartNumber`.

Type: Array of strings

Valid Values: `GetObject-Range` | `GetObject-PartNumber` | `HeadObject-Range` | `HeadObject-PartNumber`

Required: No

CloudWatchMetricsEnabled

A container for whether the CloudWatch metrics configuration is enabled.

Type: Boolean

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ObjectLambdaContentTransformation

Service: Amazon S3 Control

A container for AwsLambdaTransformation.

Contents

Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

AwsLambda

A container for an AWS Lambda function.

Type: [AwsLambdaTransformation](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ObjectLambdaTransformationConfiguration

Service: Amazon S3 Control

A configuration used when creating an Object Lambda Access Point transformation.

Contents

Actions

A container for the action of an Object Lambda Access Point configuration. Valid inputs are `GetObject`, `ListObjects`, `HeadObject`, and `ListObjectsV2`.

Type: Array of strings

Valid Values: `GetObject` | `HeadObject` | `ListObjects` | `ListObjectsV2`

Required: Yes

ContentTransformation

A container for the content transformation of an Object Lambda Access Point configuration.

Type: [ObjectLambdaContentTransformation](#) data type

Note: This object is a Union. Only one member of this object can be specified or returned.

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

PolicyStatus

Service: Amazon S3 Control

Indicates whether this access point policy is public. For more information about how Amazon S3 evaluates policies to determine whether they are public, see [The Meaning of "Public"](#) in the *Amazon S3 User Guide*.

Contents

IsPublic

Type: Boolean

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

PrefixLevel

Service: Amazon S3 Control

A container for the prefix-level configuration.

Contents

StorageMetrics

A container for the prefix-level storage metrics for S3 Storage Lens.

Type: [PrefixLevelStorageMetrics](#) data type

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

PrefixLevelStorageMetrics

Service: Amazon S3 Control

A container for the prefix-level storage metrics for S3 Storage Lens.

Contents

IsEnabled

A container for whether prefix-level storage metrics are enabled.

Type: Boolean

Required: No

SelectionCriteria

Type: [SelectionCriteria](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ProposedMultiRegionAccessPointPolicy

Service: Amazon S3 Control

The proposed access control policy for the Multi-Region Access Point.

When you update the policy, the update is first listed as the proposed policy. After the update is finished and all Regions have been updated, the proposed policy is listed as the established policy. If both policies have the same version number, the proposed policy is the established policy.

Contents

Policy

The details of the proposed policy.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

PublicAccessBlockConfiguration

Service: Amazon S3 Control

The PublicAccessBlock configuration that you want to apply to this Amazon S3 account. You can enable the configuration options in any combination. For more information about when Amazon S3 considers a bucket or object public, see [The Meaning of "Public"](#) in the *Amazon S3 User Guide*.

This data type is not supported for Amazon S3 on Outposts.

Contents

BlockPublicAcls

Specifies whether Amazon S3 should block public access control lists (ACLs) for buckets in this account. Setting this element to TRUE causes the following behavior:

- PutBucketAcl and PutObjectAcl calls fail if the specified ACL is public.
- PUT Object calls fail if the request includes a public ACL.
- PUT Bucket calls fail if the request includes a public ACL.

Enabling this setting doesn't affect existing policies or ACLs.

This property is not supported for Amazon S3 on Outposts.

Type: Boolean

Required: No

BlockPublicPolicy

Specifies whether Amazon S3 should block public bucket policies for buckets in this account. Setting this element to TRUE causes Amazon S3 to reject calls to PUT Bucket policy if the specified bucket policy allows public access.

Enabling this setting doesn't affect existing bucket policies.

This property is not supported for Amazon S3 on Outposts.

Type: Boolean

Required: No

IgnorePublicAcls

Specifies whether Amazon S3 should ignore public ACLs for buckets in this account. Setting this element to TRUE causes Amazon S3 to ignore all public ACLs on buckets in this account and any objects that they contain.

Enabling this setting doesn't affect the persistence of any existing ACLs and doesn't prevent new public ACLs from being set.

This property is not supported for Amazon S3 on Outposts.

Type: Boolean

Required: No

RestrictPublicBuckets

Specifies whether Amazon S3 should restrict public bucket policies for buckets in this account. Setting this element to TRUE restricts access to buckets with public policies to only AWS service principals and authorized users within this account.

Enabling this setting doesn't affect previously stored bucket policies, except that public and cross-account access within any public bucket policy, including non-public delegation to specific accounts, is blocked.

This property is not supported for Amazon S3 on Outposts.

Type: Boolean

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

PutMultiRegionAccessPointPolicyInput

Service: Amazon S3 Control

A container for the information associated with a [PutMultiRegionAccessPoint](#) request.

Contents

Name

The name of the Multi-Region Access Point associated with the request.

Type: String

Length Constraints: Maximum length of 50.

Pattern: `^[a-z0-9][-a-z0-9]{1,48}[a-z0-9]$`

Required: Yes

Policy

The policy details for the PutMultiRegionAccessPoint request.

Type: String

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Region

Service: Amazon S3 Control

A Region that supports a Multi-Region Access Point as well as the associated bucket for the Region.

Contents

Bucket

The name of the associated bucket for the Region.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

BucketAccountId

The AWS account ID that owns the Amazon S3 bucket that's associated with this Multi-Region Access Point.

Type: String

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

RegionalBucket

Service: Amazon S3 Control

The container for the regional bucket.

Contents

Bucket

Type: String

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: Yes

CreationDate

The creation date of the regional bucket

Type: Timestamp

Required: Yes

PublicAccessBlockEnabled

Type: Boolean

Required: Yes

BucketArn

The Amazon Resource Name (ARN) for the regional bucket.

Type: String

Length Constraints: Minimum length of 4. Maximum length of 128.

Required: No

OutpostId

The AWS Outposts ID of the regional bucket.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

RegionReport

Service: Amazon S3 Control

A combination of a bucket and Region that's part of a Multi-Region Access Point.

Contents

Bucket

The name of the bucket.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 255.

Required: No

BucketAccountId

The AWS account ID that owns the Amazon S3 bucket that's associated with this Multi-Region Access Point.

Type: String

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: No

Region

The name of the Region.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ReplicaModifications

Service: Amazon S3 Control

A filter that you can use to specify whether replica modification sync is enabled. S3 on Outposts replica modification sync can help you keep object metadata synchronized between replicas and source objects. By default, S3 on Outposts replicates metadata from the source objects to the replicas only. When replica modification sync is enabled, S3 on Outposts replicates metadata changes made to the replica copies back to the source object, making the replication bidirectional.

To replicate object metadata modifications on replicas, you can specify this element and set the Status of this element to Enabled.

Note

You must enable replica modification sync on the source and destination buckets to replicate replica metadata changes between the source and the replicas.

Contents

Status

Specifies whether S3 on Outposts replicates modifications to object metadata on replicas.

Type: String

Valid Values: Enabled | Disabled

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ReplicationConfiguration

Service: Amazon S3 Control

A container for one or more replication rules. A replication configuration must have at least one rule and you can add up to 100 rules. The maximum size of a replication configuration is 128 KB.

Contents

Role

The Amazon Resource Name (ARN) of the AWS Identity and Access Management (IAM) role that S3 on Outposts assumes when replicating objects. For information about S3 replication on Outposts configuration, see [Setting up replication](#) in the *Amazon S3 User Guide*.

Type: String

Required: Yes

Rules

A container for one or more replication rules. A replication configuration must have at least one rule and can contain an array of 100 rules at the most.

Type: Array of [ReplicationRule](#) data types

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ReplicationRule

Service: Amazon S3 Control

Specifies which S3 on Outposts objects to replicate and where to store the replicas.

Contents

Bucket

The Amazon Resource Name (ARN) of the access point for the source Outposts bucket that you want S3 on Outposts to replicate the objects from.

Type: String

Required: Yes

Destination

A container for information about the replication destination and its configurations.

Type: [Destination](#) data type

Required: Yes

Status

Specifies whether the rule is enabled.

Type: String

Valid Values: Enabled | Disabled

Required: Yes

DeleteMarkerReplication

Specifies whether S3 on Outposts replicates delete markers. If you specify a `Filter` element in your replication configuration, you must also include a `DeleteMarkerReplication` element. If your `Filter` includes a `Tag` element, the `DeleteMarkerReplication` element's `Status` child element must be set to `Disabled`, because S3 on Outposts doesn't support replicating delete markers for tag-based rules.

For more information about delete marker replication, see [How delete operations affect replication](#) in the *Amazon S3 User Guide*.

Type: [DeleteMarkerReplication](#) data type

Required: No

ExistingObjectReplication

An optional configuration to replicate existing source bucket objects.

Note

This is not supported by Amazon S3 on Outposts buckets.

Type: [ExistingObjectReplication](#) data type

Required: No

Filter

A filter that identifies the subset of objects to which the replication rule applies. A `Filter` element must specify exactly one `Prefix`, `Tag`, or `And` child element.

Type: [ReplicationRuleFilter](#) data type

Required: No

ID

A unique identifier for the rule. The maximum value is 255 characters.

Type: String

Required: No

Prefix

This member has been deprecated.

An object key name prefix that identifies the object or objects to which the rule applies. The maximum prefix length is 1,024 characters. To include all objects in an Outposts bucket, specify an empty string.

Important

When you're using XML requests, you must replace special characters (such as carriage returns) in object keys with their equivalent XML entity codes. For more information, see [XML-related object key constraints](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

Priority

The priority indicates which rule has precedence whenever two or more replication rules conflict. S3 on Outposts attempts to replicate objects according to all replication rules. However, if there are two or more rules with the same destination Outposts bucket, then objects will be replicated according to the rule with the highest priority. The higher the number, the higher the priority.

For more information, see [Creating replication rules on Outposts](#) in the *Amazon S3 User Guide*.

Type: Integer

Required: No

SourceSelectionCriteria

A container that describes additional filters for identifying the source Outposts objects that you want to replicate. You can choose to enable or disable the replication of these objects.

Type: [SourceSelectionCriteria](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ReplicationRuleAndOperator

Service: Amazon S3 Control

A container for specifying rule filters. The filters determine the subset of objects to which the rule applies. This element is required only if you specify more than one filter.

For example:

- If you specify both a `Prefix` and a `Tag` filter, wrap these filters in an `And` element.
- If you specify a filter based on multiple tags, wrap the `Tag` elements in an `And` element.

Contents

Prefix

An object key name prefix that identifies the subset of objects that the rule applies to.

Type: String

Required: No

Tags

An array of tags that contain key and value pairs.

Type: Array of [S3Tag](#) data types

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ReplicationRuleFilter

Service: Amazon S3 Control

A filter that identifies the subset of objects to which the replication rule applies. A `Filter` element must specify exactly one `Prefix`, `Tag`, or `And` child element.

Contents

And

A container for specifying rule filters. The filters determine the subset of objects that the rule applies to. This element is required only if you specify more than one filter. For example:

- If you specify both a `Prefix` and a `Tag` filter, wrap these filters in an `And` element.
- If you specify a filter based on multiple tags, wrap the `Tag` elements in an `And` element.

Type: [ReplicationRuleAndOperator](#) data type

Required: No

Prefix

An object key name prefix that identifies the subset of objects that the rule applies to.

Important

When you're using XML requests, you must replace special characters (such as carriage returns) in object keys with their equivalent XML entity codes. For more information, see [XML-related object key constraints](#) in the *Amazon S3 User Guide*.

Type: String

Required: No

Tag

A container for a key-value name pair.

Type: [S3Tag](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ReplicationTime

Service: Amazon S3 Control

A container that specifies S3 Replication Time Control (S3 RTC) related information, including whether S3 RTC is enabled and the time when all objects and operations on objects must be replicated.

Note

This is not supported by Amazon S3 on Outposts buckets.

Contents

Status

Specifies whether S3 Replication Time Control (S3 RTC) is enabled.

Type: String

Valid Values: Enabled | Disabled

Required: Yes

Time

A container that specifies the time by which replication should be complete for all objects and operations on objects.

Type: [ReplicationTimeValue](#) data type

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)

- [AWS SDK for Ruby V3](#)

ReplicationTimeValue

Service: Amazon S3 Control

A container that specifies the time value for S3 Replication Time Control (S3 RTC). This value is also used for the replication metrics EventThreshold element.

Note

This is not supported by Amazon S3 on Outposts buckets.

Contents

Minutes

Contains an integer that specifies the time period in minutes.

Valid value: 15

Type: Integer

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3AccessControlList

Service: Amazon S3 Control

Contents

Owner

Type: [S3ObjectOwner](#) data type

Required: Yes

Grants

Type: Array of [S3Grant](#) data types

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3AccessControlPolicy

Service: Amazon S3 Control

Contents

AccessControlList

Type: [S3AccessControlList](#) data type

Required: No

CannedAccessControlList

Type: String

Valid Values: `private` | `public-read` | `public-read-write` | `aws-exec-read` | `authenticated-read` | `bucket-owner-read` | `bucket-owner-full-control`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3BucketDestination

Service: Amazon S3 Control

A container for the bucket where the Amazon S3 Storage Lens metrics export files are located.

Contents

AccountId

The account ID of the owner of the S3 Storage Lens metrics export bucket.

Type: String

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: Yes

Arn

The Amazon Resource Name (ARN) of the bucket. This property is read-only and follows the following format: `arn:aws:s3:us-east-1:example-account-id:bucket/your-destination-bucket-name`

Type: String

Length Constraints: Minimum length of 1. Maximum length of 128.

Pattern: `arn:[^:]+:s3:.*`

Required: Yes

Format

Type: String

Valid Values: CSV | Parquet

Required: Yes

OutputSchemaVersion

The schema version of the export file.

Type: String

Valid Values: V_1

Required: Yes

Encryption

The container for the type encryption of the metrics exports in this bucket.

Type: [StorageLensDataExportEncryption](#) data type

Required: No

Prefix

The prefix of the destination bucket where the metrics export will be delivered.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3ComputeObjectChecksumOperation

Service: Amazon S3 Control

Directs the specified job to invoke the ComputeObjectChecksum operation on every object listed in the job's manifest.

Contents

ChecksumAlgorithm

Indicates the algorithm that you want Amazon S3 to use to create the checksum. For more information, see [Checking object integrity](#) in the Amazon S3 User Guide.

Type: String

Valid Values: CRC32 | CRC32C | CRC64NVME | MD5 | SHA1 | SHA256 | SHA512 | XXHASH64 | XXHASH3 | XXHASH128

Required: No

ChecksumType

Indicates the checksum type that you want Amazon S3 to use to calculate the object's checksum value. For more information, see [Checking object integrity](#) in the Amazon S3 User Guide.

Type: String

Valid Values: FULL_OBJECT | COMPOSITE

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3CopyObjectOperation

Service: Amazon S3 Control

Contains the configuration parameters for a PUT Copy object operation. S3 Batch Operations passes every object to the underlying CopyObject API operation. For more information about the parameters for this operation, see [CopyObject](#).

Contents

AccessControlGrants

Note

This functionality is not supported by directory buckets.

Type: Array of [S3Grant](#) data types

Required: No

AnnotationDirective

Specifies whether the Batch Operations copy job copies object annotations from the source object or skips them. If this property isn't specified, COPY is the default behavior.

Valid Values: COPY | EXCLUDE

Note

This functionality is not supported by directory buckets.

Type: String

Valid Values: COPY | EXCLUDE

Required: No

BucketKeyEnabled

Specifies whether Amazon S3 should use an S3 Bucket Key for object encryption with server-side encryption using AWS KMS (SSE-KMS). Setting this header to `true` causes Amazon S3 to use an S3 Bucket Key for object encryption with SSE-KMS.

Specifying this header with an *Copy* action doesn't affect *bucket-level* settings for S3 Bucket Key.

Note

Directory buckets - S3 Bucket Keys aren't supported, when you copy SSE-KMS encrypted objects from general purpose buckets to directory buckets, from directory buckets to general purpose buckets, or between directory buckets, through [the Copy operation in Batch Operations](#). In this case, Amazon S3 makes a call to AWS KMS every time a copy request is made for a KMS-encrypted object.

Type: Boolean

Required: No

CannedAccessControlList

Note

This functionality is not supported by directory buckets.

Type: String

Valid Values: `private` | `public-read` | `public-read-write` | `aws-exec-read` | `authenticated-read` | `bucket-owner-read` | `bucket-owner-full-control`

Required: No

ChecksumAlgorithm

Indicates the algorithm that you want Amazon S3 to use to create the checksum. For more information, see [Checking object integrity](#) in the *Amazon S3 User Guide*.

Type: String

Valid Values: CRC32 | CRC32C | SHA1 | SHA256 | CRC64NVME | SHA512 | MD5 |
XXHASH64 | XXHASH3 | XXHASH128

Required: No

MetadataDirective

Type: String

Valid Values: COPY | REPLACE

Required: No

ModifiedSinceConstraint

Type: Timestamp

Required: No

NewObjectMetadata

If you don't provide this parameter, Amazon S3 copies all the metadata from the original objects. If you specify an empty set, the new objects will have no tags. Otherwise, Amazon S3 assigns the supplied tags to the new objects.

Type: [S3ObjectMetadata](#) data type

Required: No

NewObjectTagging

Specifies a list of tags to add to the destination objects after they are copied. If NewObjectTagging is not specified, the tags of the source objects are copied to destination objects by default.

Note

Directory buckets - Tags aren't supported by directory buckets. If your source objects have tags and your destination bucket is a directory bucket, specify an empty tag set in the NewObjectTagging field to prevent copying the source object tags to the directory bucket.

Type: Array of [S3Tag](#) data types

Required: No

ObjectLockEventHold

The event hold status to be applied to all objects in the Batch Operations copy job. Set to ON to enable an event hold or OFF to disable it.

Note

This functionality is not supported by directory buckets.

Type: String

Valid Values: ON | OFF

Required: No

ObjectLockEventHoldDuration

The event hold duration to be applied to all objects in the Batch Operations copy job. The duration specifies how long the object remains protected after the event hold is released.

Note

This functionality is not supported by directory buckets.

Type: [S3ObjectLockEventHoldDuration](#) data type

Required: No

ObjectLockLegalHoldStatus

The legal hold status to be applied to all objects in the Batch Operations job.

Note

This functionality is not supported by directory buckets.

Type: String

Valid Values: OFF | ON

Required: No

ObjectLockMode

The retention mode to be applied to all objects in the Batch Operations job.

Note

This functionality is not supported by directory buckets.

Type: String

Valid Values: COMPLIANCE | GOVERNANCE

Required: No

ObjectLockRetainUntilDate

The date when the applied object retention configuration expires on all objects in the Batch Operations job.

Note

This functionality is not supported by directory buckets.

Type: Timestamp

Required: No

RedirectLocation

If the destination bucket is configured as a website, specifies an optional metadata property for website redirects, `x-amz-website-redirect-location`. Allows webpage redirects if the object copy is accessed through a website endpoint.

Note

This functionality is not supported by directory buckets.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: No

RequesterPays

Note

This functionality is not supported by directory buckets.

Type: Boolean

Required: No

SSEAwsKmsKeyId

Specifies the AWS KMS key ID (Key ID, Key ARN, or Key Alias) to use for object encryption. If the KMS key doesn't exist in the same account that's issuing the command, you must use the full Key ARN not the Key ID.

Note

Directory buckets - If you specify `SSEAlgorithm` with KMS, you must specify the `SSEAwsKmsKeyId` parameter with the ID (Key ID or Key ARN) of the AWS KMS symmetric encryption customer managed key to use. Otherwise, you get an HTTP 400 Bad Request error. The key alias format of the KMS key isn't supported. To encrypt new object copies in a directory bucket with SSE-KMS, you must specify SSE-KMS as the directory bucket's default encryption configuration with a KMS key (specifically, a [customer managed key](#)). The [AWS managed key](#) (aws/s3) isn't supported. Your SSE-KMS configuration can only support 1 [customer managed key](#) per directory bucket for the lifetime of the bucket. After you specify a customer managed key for SSE-KMS as the bucket default encryption, you can't override the customer managed key for the bucket's SSE-KMS configuration. Then, when you specify server-side encryption settings for new object copies with SSE-KMS, you must make sure the encryption key is the same customer managed key that you specified for the directory bucket's default encryption configuration.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2000.

Required: No

StorageClass

Specify the storage class for the destination objects in a Copy operation.

Note

Directory buckets - This functionality is not supported by directory buckets.

Type: String

Valid Values: STANDARD | STANDARD_IA | ONEZONE_IA | GLACIER | INTELLIGENT_TIERING | DEEP_ARCHIVE | GLACIER_IR

Required: No

TargetKeyPrefix

Specifies the folder prefix that you want the objects to be copied into. For example, to copy objects into a folder named `Folder1` in the destination bucket, set the `TargetKeyPrefix` property to `Folder1`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

TargetResource

Specifies the destination bucket Amazon Resource Name (ARN) for the batch copy operation.

- **General purpose buckets** - For example, to copy objects to a general purpose bucket named `destinationBucket`, set the `TargetResource` property to `arn:aws:s3:::destinationBucket`.
- **Directory buckets** - For example, to copy objects to a directory bucket named `destinationBucket` in the Availability Zone identified by the AZ ID `usw2-az1`, set

the `TargetResource` property to `arn:aws:s3express:region:account_id:/bucket/destination_bucket_base_name--usw2-az1--x-s3`. A directory bucket as a destination bucket can be in Availability Zone or Local Zone.

 Note

Copying objects across different AWS Regions isn't supported when the source or destination bucket is in AWS Local Zones. The source and destination buckets must have the same parent AWS Region. Otherwise, you get an HTTP 400 Bad Request error with the error code `InvalidRequest`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 128.

Pattern: `arn:[^:]+:(s3|s3express):.*`

Required: No

UnModifiedSinceConstraint

Type: Timestamp

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3DeleteObjectTaggingOperation

Service: Amazon S3 Control

Contains no configuration parameters because the DELETE Object tagging (DeleteObjectTagging) API operation accepts only the bucket name and key name as parameters, which are defined in the job's manifest.

Contents

The members of this exception structure are context-dependent.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3GeneratedManifestDescriptor

Service: Amazon S3 Control

Describes the specified job's generated manifest. Batch Operations jobs created with a ManifestGenerator populate details of this descriptor after execution of the ManifestGenerator.

Contents

Format

The format of the generated manifest.

Type: String

Valid Values: S3InventoryReport_CSV_20211130

Required: No

Location

Contains the information required to locate a manifest object. Manifests can't be imported from directory buckets. For more information, see [Directory buckets](#).

Type: [JobManifestLocation](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3Grant

Service: Amazon S3 Control

Contents

Grantee

Type: [S3Grantee](#) data type

Required: No

Permission

Type: String

Valid Values: FULL_CONTROL | READ | WRITE | READ_ACP | WRITE_ACP

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3Grantee

Service: Amazon S3 Control

Contents

DisplayName

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

Identifier

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

TypeIdentifier

Type: String

Valid Values: `id` | `emailAddress` | `uri`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3InitiateRestoreObjectOperation

Service: Amazon S3 Control

Contains the configuration parameters for a POST Object restore job. S3 Batch Operations passes every object to the underlying RestoreObject API operation. For more information about the parameters for this operation, see [RestoreObject](#).

Contents

ExpirationInDays

This argument specifies how long the S3 Glacier or S3 Glacier Deep Archive object remains available in Amazon S3. S3 Initiate Restore Object jobs that target S3 Glacier and S3 Glacier Deep Archive objects require ExpirationInDays set to 1 or greater.

Conversely, do *not* set ExpirationInDays when creating S3 Initiate Restore Object jobs that target S3 Intelligent-Tiering Archive Access and Deep Archive Access tier objects. Objects in S3 Intelligent-Tiering archive access tiers are not subject to restore expiry, so specifying ExpirationInDays results in restore request failure.

S3 Batch Operations jobs can operate either on S3 Glacier and S3 Glacier Deep Archive storage class objects or on S3 Intelligent-Tiering Archive Access and Deep Archive Access storage tier objects, but not both types in the same job. If you need to restore objects of both types you *must* create separate Batch Operations jobs.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

GlacierJobTier

S3 Batch Operations supports STANDARD and BULK retrieval tiers, but not the EXPEDITED retrieval tier.

Type: String

Valid Values: BULK | STANDARD

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3JobManifestGenerator

Service: Amazon S3 Control

The container for the service that will create the S3 manifest.

Contents

EnableManifestOutput

Determines whether or not to write the job's generated manifest to a bucket.

Type: Boolean

Required: Yes

SourceBucket

The ARN of the source bucket used by the ManifestGenerator.

Note

Directory buckets - Directory buckets aren't supported as the source buckets used by S3JobManifestGenerator to generate the job manifest.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 128.

Pattern: `arn:[^:]+:s3:.*`

Required: Yes

ExpectedBucketOwner

The AWS account ID that owns the source bucket specified in SourceBucket. If provided, the manifest source bucket owner's AWS account ID must match this value, else the job fails.

Type: String

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: No

Filter

Specifies rules the S3JobManifestGenerator should use to decide whether an object in the source bucket should or should not be included in the generated job manifest.

Type: [JobManifestGeneratorFilter](#) data type

Required: No

ManifestOutputLocation

Specifies the location the generated manifest will be written to. Manifests can't be written to directory buckets. For more information, see [Directory buckets](#).

Type: [S3ManifestOutputLocation](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3ManifestOutputLocation

Service: Amazon S3 Control

Location details for where the generated manifest should be written.

Contents

Bucket

The bucket ARN the generated manifest should be written to.

Note

Directory buckets - Directory buckets aren't supported as the buckets to store the generated manifest.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 128.

Pattern: `arn:[^:]+:s3:.*`

Required: Yes

ManifestFormat

The format of the generated manifest.

Type: String

Valid Values: `S3InventoryReport_CSV_20211130`

Required: Yes

ExpectedManifestBucketOwner

The Account ID that owns the bucket the generated manifest is written to.

Type: String

Length Constraints: Maximum length of 64.

Pattern: `^\d{12}$`

Required: No

ManifestEncryption

Specifies what encryption should be used when the generated manifest objects are written.

Type: [GeneratedManifestEncryption](#) data type

Required: No

ManifestPrefix

Prefix identifying one or more objects to which the manifest applies.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3ObjectLockEventHoldDuration

Service: Amazon S3 Control

Contains the duration configuration for an event hold, specified in either days or years.

Contents

Days

The number of days for the event hold duration. The minimum value is 1 and the maximum value is 36,500.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 36500.

Required: No

Years

The number of years for the event hold duration. The minimum value is 1 and the maximum value is 100.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 100.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3ObjectLockLegalHold

Service: Amazon S3 Control

Whether S3 Object Lock legal hold will be applied to objects in an S3 Batch Operations job.

Contents

Status

The Object Lock legal hold status to be applied to all objects in the Batch Operations job.

Type: String

Valid Values: OFF | ON

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3ObjectLockRetentionEventHoldDuration

Service: Amazon S3 Control

Contains the duration configuration for an event hold, specified in either days or years.

Contents

Days

The number of days for the event hold duration. The minimum value is 1 and the maximum value is 36,500.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 36500.

Required: No

Years

The number of years for the event hold duration. The minimum value is 1 and the maximum value is 100.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 100.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3ObjectMetadata

Service: Amazon S3 Control

Contents

CacheControl

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

ContentDisposition

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

ContentEncoding

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

ContentLanguage

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

ContentLength

This member has been deprecated.

Type: Long

Valid Range: Minimum value of 0.

Required: No

ContentMD5

This member has been deprecated.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

ContentType

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

HttpExpiresDate

Type: Timestamp

Required: No

RequesterCharged

This member has been deprecated.

Type: Boolean

Required: No

SSEAlgorithm

The server-side encryption algorithm used when storing objects in Amazon S3.

Directory buckets - For directory buckets, there are only two supported options for server-side encryption: server-side encryption with Amazon S3 managed keys (SSE-S3) (AES256)

and server-side encryption with AWS KMS keys (SSE-KMS) (KMS). For more information, see [Protecting data with server-side encryption](#) in the *Amazon S3 User Guide*. For [the Copy operation in Batch Operations](#), see [S3CopyObjectOperation](#).

Type: String

Valid Values: AES256 | KMS

Required: No

UserMetadata

Type: String to string map

Map Entries: Maximum number of 8192 items.

Key Length Constraints: Minimum length of 1. Maximum length of 1024.

Value Length Constraints: Maximum length of 1024.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3ObjectOwner

Service: Amazon S3 Control

Contents

DisplayName

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

ID

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3ReplicateObjectOperation

Service: Amazon S3 Control

Directs the specified job to invoke `ReplicateObject` on every object in the job's manifest.

Contents

The members of this exception structure are context-dependent.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3Retention

Service: Amazon S3 Control

Contains the S3 Object Lock retention mode to be applied to all objects in the S3 Batch Operations job. If you don't provide Mode and RetainUntilDate data types in your operation, you will remove the retention from your objects. For more information, see [Using S3 Object Lock retention with S3 Batch Operations](#) in the *Amazon S3 User Guide*.

Contents

EventHold

The event hold status to be applied to all objects in the Batch Operations job. Set to ON to enable an event hold or OFF to disable it.

Type: String

Valid Values: ON | OFF

Required: No

EventHoldDuration

The event hold duration to be applied to all objects in the Batch Operations job. The duration specifies how long the object remains protected after the event hold is released.

Type: [S3ObjectLockRetentionEventHoldDuration](#) data type

Required: No

Mode

The Object Lock retention mode to be applied to all objects in the Batch Operations job.

Type: String

Valid Values: COMPLIANCE | GOVERNANCE

Required: No

RetainUntilDate

The date when the applied Object Lock retention will expire on all objects set by the Batch Operations job.

Type: Timestamp

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3SetObjectAclOperation

Service: Amazon S3 Control

Contains the configuration parameters for a PUT Object ACL operation. S3 Batch Operations passes every object to the underlying `PutObjectAcl` API operation. For more information about the parameters for this operation, see [PutObjectAcl](#).

Contents

AccessControlPolicy

Type: [S3AccessControlPolicy](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3SetObjectLegalHoldOperation

Service: Amazon S3 Control

Contains the configuration for an S3 Object Lock legal hold operation that an S3 Batch Operations job passes to every object to the underlying PutObjectLegalHold API operation. For more information, see [Using S3 Object Lock legal hold with S3 Batch Operations](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported by directory buckets.

Contents

LegalHold

Contains the Object Lock legal hold status to be applied to all objects in the Batch Operations job.

Type: [S3ObjectLockLegalHold](#) data type

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3SetObjectRetentionOperation

Service: Amazon S3 Control

Contains the configuration parameters for the Object Lock retention action for an S3 Batch Operations job. Batch Operations passes every object to the underlying `PutObjectRetention` API operation. For more information, see [Using S3 Object Lock retention with S3 Batch Operations](#) in the *Amazon S3 User Guide*.

Note

This functionality is not supported by directory buckets.

Contents

Retention

Contains the Object Lock retention mode to be applied to all objects in the Batch Operations job. For more information, see [Using S3 Object Lock retention with S3 Batch Operations](#) in the *Amazon S3 User Guide*.

Type: [S3Retention](#) data type

Required: Yes

BypassGovernanceRetention

Indicates if the action should be applied to objects in the Batch Operations job even if they have Object Lock GOVERNANCE type in place.

Type: Boolean

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3SetObjectTaggingOperation

Service: Amazon S3 Control

Contains the configuration parameters for a PUT Object Tagging operation. S3 Batch Operations passes every object to the underlying PutObjectTagging API operation. For more information about the parameters for this operation, see [PutObjectTagging](#).

Contents

TagSet

Type: Array of [S3Tag](#) data types

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3Tag

Service: Amazon S3 Control

A container for a key-value name pair.

Contents

Key

Key of the tag

Type: String

Length Constraints: Minimum length of 1. Maximum length of 128.

Pattern: `^([\p{L}\p{Z}\p{N}_.: /+=\ -@]*)$`

Required: Yes

Value

Value of the tag

Type: String

Length Constraints: Minimum length of 0. Maximum length of 256.

Pattern: `^([\p{L}\p{Z}\p{N}_.: /+=\ -@]*)$`

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3UpdateObjectEncryptionOperation

Service: Amazon S3 Control

With the UpdateObjectEncryption operation, you can atomically update the server-side encryption type of an existing object in a general purpose bucket without any data movement.

Contents

ObjectEncryption

The updated server-side encryption type for this S3 object. The UpdateObjectEncryption operation supports the SSE-KMS encryption type.

Type: [ObjectEncryption](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3UpdateObjectEncryptionSSEKMS

Service: Amazon S3 Control

If SSEKMS is specified for UpdateObjectEncryption, this data type specifies the AWS KMS key Amazon Resource Name (ARN) to use and whether to use an S3 Bucket Key for server-side encryption using AWS Key Management Service (AWS KMS) keys (SSE-KMS).

Contents

KMSKeyArn

Specifies the AWS KMS key Amazon Resource Name (ARN) to use for the updated server-side encryption type. Required if UpdateObjectEncryption specifies SSEKMS.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `arn:aws[a-zA-Z0-9-]*:kms:[a-z0-9-]+:[0-9]{12}:key/[a-zA-Z0-9-]+`

Required: Yes

BucketKeyEnabled

Specifies whether Amazon S3 should use an S3 Bucket Key for object encryption with server-side encryption using AWS Key Management Service (AWS KMS) keys (SSE-KMS). If this value isn't specified, it defaults to `false`. Setting this value to `true` causes Amazon S3 to use an S3 Bucket Key for update object encryption with SSE-KMS.

Type: Boolean

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Scope

Service: Amazon S3 Control

You can use the access point scope to restrict access to specific prefixes, API operations, or a combination of both.

For more information, see [Manage the scope of your access points for directory buckets](#).

Contents

Permissions

You can include one or more API operations as permissions.

Type: Array of strings

Valid Values: `GetObject` | `GetObjectAttributes` | `ListMultipartUploadParts` | `ListBucket` | `ListBucketMultipartUploads` | `PutObject` | `DeleteObject` | `AbortMultipartUpload`

Required: No

Prefixes

You can specify any amount of prefixes, but the total length of characters of all prefixes must be less than 256 bytes in size.

Type: Array of strings

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SelectionCriteria

Service: Amazon S3 Control

Contents

Delimiter

A container for the delimiter of the selection criteria being used.

Type: String

Length Constraints: Maximum length of 1.

Required: No

MaxDepth

The max depth of the selection criteria

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 10.

Required: No

MinStorageBytesPercentage

The minimum percentage of total bucket storage that a prefix must hold for its metrics to be included.

Type: Double

Valid Range: Minimum value of 1. Maximum value of 100.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SourceSelectionCriteria

Service: Amazon S3 Control

A container that describes additional filters for identifying the source objects that you want to replicate. You can choose to enable or disable the replication of these objects.

Contents

ReplicaModifications

A filter that you can use to specify whether replica modification sync is enabled. S3 on Outposts replica modification sync can help you keep object metadata synchronized between replicas and source objects. By default, S3 on Outposts replicates metadata from the source objects to the replicas only. When replica modification sync is enabled, S3 on Outposts replicates metadata changes made to the replica copies back to the source object, making the replication bidirectional.

To replicate object metadata modifications on replicas, you can specify this element and set the Status of this element to Enabled.

Note

You must enable replica modification sync on the source and destination buckets to replicate replica metadata changes between the source and the replicas.

Type: [ReplicaModifications](#) data type

Required: No

SseKmsEncryptedObjects

A filter that you can use to select Amazon S3 objects that are encrypted with server-side encryption by using AWS Key Management Service (AWS KMS) keys. If you include SourceSelectionCriteria in the replication configuration, this element is required.

Note

This is not supported by Amazon S3 on Outposts buckets.

Type: [SseKmsEncryptedObjects](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SSECFilter

Service: Amazon S3 Control

A filter that returns objects that are encrypted by server-side encryption with customer-provided keys (SSE-C).

Contents

The members of this exception structure are context-dependent.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SSEKMS

Service: Amazon S3 Control

Contents

KeyId

A container for the ARN of the SSE-KMS encryption. This property is read-only and follows the following format: `arn:aws:kms:us-east-1:example-account-id:key/example-9a73-4afc-8d29-8f5900cef44e`

Type: String

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SseKmsEncryptedObjects

Service: Amazon S3 Control

A container for filter information that you can use to select S3 objects that are encrypted with AWS Key Management Service (AWS KMS).

Note

This is not supported by Amazon S3 on Outposts buckets.

Contents

Status

Specifies whether Amazon S3 replicates objects that are created with server-side encryption by using an AWS KMS key stored in AWS Key Management Service.

Type: String

Valid Values: Enabled | Disabled

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SSEKMSEncryption

Service: Amazon S3 Control

Configuration for the use of SSE-KMS to encrypt generated manifest objects.

Contents

KeyId

Specifies the ID of the AWS Key Management Service (AWS KMS) symmetric encryption customer managed key to use for encrypting generated manifest objects.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2000.

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SSEKMSFilter

Service: Amazon S3 Control

A filter that returns objects that are encrypted by server-side encryption with AWS KMS (SSE-KMS).

Contents

BucketKeyEnabled

Specifies whether Amazon S3 should use an S3 Bucket Key for object encryption with server-side encryption using AWS Key Management Service (AWS KMS) keys (SSE-KMS). If specified, will filter SSE-KMS encrypted objects by S3 Bucket Key status.

Type: Boolean

Required: No

KmsKeyArn

The Amazon Resource Name (ARN) of the customer managed KMS key to use for the filter to return objects that are encrypted by the specified key. For best performance, use keys in the same Region as the S3 Batch Operations job.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `arn:aws[a-zA-Z0-9-]*:kms:[a-z0-9-]+:[0-9]{12}:key/[a-zA-Z0-9-]+`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SSES3

Service: Amazon S3 Control

Contents

The members of this exception structure are context-dependent.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SSE3Encryption

Service: Amazon S3 Control

Configuration for the use of SSE-S3 to encrypt generated manifest objects.

Contents

The members of this exception structure are context-dependent.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SSES3Filter

Service: Amazon S3 Control

A filter that returns objects that are encrypted by server-side encryption with Amazon S3 managed keys (SSE-S3).

Contents

The members of this exception structure are context-dependent.

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

StorageLensAwsOrg

Service: Amazon S3 Control

The AWS organization for your S3 Storage Lens.

Contents

Arn

A container for the Amazon Resource Name (ARN) of the AWS organization. This property is read-only and follows the following format: `arn:aws:organizations:us-east-1:example-account-id:organization/o-ex2l495dck`

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `arn:[a-z\-\+]:organizations::\d{12}:organization/o-[a-z0-9]{10,32}`

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

StorageLensConfiguration

Service: Amazon S3 Control

A container for the Amazon S3 Storage Lens configuration.

Contents

AccountLevel

A container for all the account-level configurations of your S3 Storage Lens configuration.

Type: [AccountLevel](#) data type

Required: Yes

Id

A container for the Amazon S3 Storage Lens configuration ID.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\_\.\.]+`

Required: Yes

IsEnabled

A container for whether the S3 Storage Lens configuration is enabled.

Type: Boolean

Required: Yes

AwsOrg

A container for the AWS organization for this S3 Storage Lens configuration.

Type: [StorageLensAwsOrg](#) data type

Required: No

DataExport

A container to specify the properties of your S3 Storage Lens metrics export including, the destination, schema and format.

Type: [StorageLensDataExport](#) data type

Required: No

Exclude

A container for what is excluded in this configuration. This container can only be valid if there is no `Include` container submitted, and it's not empty.

Type: [Exclude](#) data type

Required: No

ExpandedPrefixesDataExport

A container that configures your S3 Storage Lens expanded prefixes metrics report.

Type: [StorageLensExpandedPrefixesDataExport](#) data type

Required: No

Include

A container for what is included in this configuration. This container can only be valid if there is no `Exclude` container submitted, and it's not empty.

Type: [Include](#) data type

Required: No

PrefixDelimiter

A container for all prefix delimiters that are used for object keys in this S3 Storage Lens configuration. The prefix delimiters determine how S3 Storage Lens counts prefix depth, by separating the hierarchical levels in object keys.

Note

- If either a prefix delimiter or existing delimiter is undefined, Amazon S3 uses the delimiter that's defined.
- If both the prefix delimiter and existing delimiter are undefined, S3 uses `/` as the default delimiter.
- When custom delimiters are used, both the prefix delimiter and existing delimiter must specify the same special character. Otherwise, your request results in an error.

Type: String

Length Constraints: Maximum length of 1.

Required: No

StorageLensArn

The Amazon Resource Name (ARN) of the S3 Storage Lens configuration. This property is read-only and follows the following format: `arn:aws:s3:us-east-1:example-account-id:storage-lens/your-dashboard-name`

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `arn:[a-z\-\+]:s3:[a-z0-9\-\+]:\d{12}:storage\-\lens\/.*`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

StorageLensDataExport

Service: Amazon S3 Control

A container to specify the properties of your S3 Storage Lens metrics export, including the destination, schema, and format.

Contents

CloudWatchMetrics

A container for enabling Amazon CloudWatch publishing for S3 Storage Lens metrics.

Type: [CloudWatchMetrics](#) data type

Required: No

S3BucketDestination

A container for the bucket where the S3 Storage Lens metrics export will be located.

Note

This bucket must be located in the same Region as the storage lens configuration.

Type: [S3BucketDestination](#) data type

Required: No

StorageLensTableDestination

A container for configuring S3 Storage Lens data exports to read-only S3 table buckets.

Type: [StorageLensTableDestination](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

StorageLensDataExportEncryption

Service: Amazon S3 Control

A container for the encryption of the S3 Storage Lens metrics exports.

Contents

SSEKMS

Type: [SSEKMS](#) data type

Required: No

SSES3

Type: [SSES3](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

StorageLensExpandedPrefixesDataExport

Service: Amazon S3 Control

A container for your S3 Storage Lens expanded prefix metrics report configuration. Unlike the default Storage Lens metrics report, the enhanced prefix metrics report includes all S3 Storage Lens storage and activity data related to the full list of prefixes in your Storage Lens configuration.

Contents

S3BucketDestination

A container for the bucket where the Amazon S3 Storage Lens metrics export files are located.

Type: [S3BucketDestination](#) data type

Required: No

StorageLensTableDestination

A container for the bucket where the S3 Storage Lens metric export files are located. At least one export destination must be specified.

Type: [StorageLensTableDestination](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

StorageLensGroup

Service: Amazon S3 Control

A custom grouping of objects that include filters for prefixes, suffixes, object tags, object size, or object age. You can create an S3 Storage Lens group that includes a single filter or multiple filter conditions. To specify multiple filter conditions, you use AND or OR logical operators.

Contents

Filter

Sets the criteria for the Storage Lens group data that is displayed. For multiple filter conditions, the AND or OR logical operator is used.

Type: [StorageLensGroupFilter](#) data type

Required: Yes

Name

Contains the name of the Storage Lens group.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[a-zA-Z0-9\-\_\+]`

Required: Yes

StorageLensGroupArn

Contains the Amazon Resource Name (ARN) of the Storage Lens group. This property is read-only.

Type: String

Length Constraints: Minimum length of 4. Maximum length of 1024.

Pattern: `arn:[a-z\-\+]:s3:[a-z0-9\-\+]:\d{12}:storage\-lens\-group\/.*`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

StorageLensGroupAndOperator

Service: Amazon S3 Control

A logical operator that allows multiple filter conditions to be joined for more complex comparisons of Storage Lens group data.

Contents

MatchAnyPrefix

Contains a list of prefixes. At least one prefix must be specified. Up to 10 prefixes are allowed.

Type: Array of strings

Required: No

MatchAnySuffix

Contains a list of suffixes. At least one suffix must be specified. Up to 10 suffixes are allowed.

Type: Array of strings

Required: No

MatchAnyTag

Contains the list of object tags. At least one object tag must be specified. Up to 10 object tags are allowed.

Type: Array of [S3Tag](#) data types

Required: No

MatchObjectAge

Contains DaysGreaterThan and DaysLessThan to define the object age range (minimum and maximum number of days).

Type: [MatchObjectAge](#) data type

Required: No

MatchObjectSize

Contains BytesGreaterThan and BytesLessThan to define the object size range (minimum and maximum number of Bytes).

Type: [MatchObjectSize](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

StorageLensGroupFilter

Service: Amazon S3 Control

The filter element sets the criteria for the Storage Lens group data that is displayed. For multiple filter conditions, the AND or OR logical operator is used.

Contents

And

A logical operator that allows multiple filter conditions to be joined for more complex comparisons of Storage Lens group data. Objects must match all of the listed filter conditions that are joined by the And logical operator. Only one of each filter condition is allowed.

Type: [StorageLensGroupAndOperator](#) data type

Required: No

MatchAnyPrefix

Contains a list of prefixes. At least one prefix must be specified. Up to 10 prefixes are allowed.

Type: Array of strings

Required: No

MatchAnySuffix

Contains a list of suffixes. At least one suffix must be specified. Up to 10 suffixes are allowed.

Type: Array of strings

Required: No

MatchAnyTag

Contains the list of S3 object tags. At least one object tag must be specified. Up to 10 object tags are allowed.

Type: Array of [S3Tag](#) data types

Required: No

MatchObjectAge

Contains `DaysGreaterThan` and `DaysLessThan` to define the object age range (minimum and maximum number of days).

Type: [MatchObjectAge](#) data type

Required: No

MatchObjectSize

Contains `BytesGreaterThan` and `BytesLessThan` to define the object size range (minimum and maximum number of Bytes).

Type: [MatchObjectSize](#) data type

Required: No

Or

A single logical operator that allows multiple filter conditions to be joined. Objects can match any of the listed filter conditions, which are joined by the `Or` logical operator. Only one of each filter condition is allowed.

Type: [StorageLensGroupOrOperator](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

StorageLensGroupLevel

Service: Amazon S3 Control

Specifies the Storage Lens groups to include in the Storage Lens group aggregation.

Contents

SelectionCriteria

Indicates which Storage Lens group ARNs to include or exclude in the Storage Lens group aggregation. If this value is left null, then all Storage Lens groups are selected.

Type: [StorageLensGroupLevelSelectionCriteria](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

StorageLensGroupLevelSelectionCriteria

Service: Amazon S3 Control

Indicates which Storage Lens group ARNs to include or exclude in the Storage Lens group aggregation. You can only attach Storage Lens groups to your Storage Lens dashboard if they're included in your Storage Lens group aggregation. If this value is left null, then all Storage Lens groups are selected.

Contents

Exclude

Indicates which Storage Lens group ARNs to exclude from the Storage Lens group aggregation.

Type: Array of strings

Length Constraints: Minimum length of 4. Maximum length of 1024.

Pattern: `arn:[a-z\-\-]+:s3:[a-z0-9\-\-]+:\d{12}:storage\-\lens\-\group\/.*`

Required: No

Include

Indicates which Storage Lens group ARNs to include in the Storage Lens group aggregation.

Type: Array of strings

Length Constraints: Minimum length of 4. Maximum length of 1024.

Pattern: `arn:[a-z\-\-]+:s3:[a-z0-9\-\-]+:\d{12}:storage\-\lens\-\group\/.*`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)

- [AWS SDK for Ruby V3](#)

StorageLensGroupOrOperator

Service: Amazon S3 Control

A container element for specifying `Or` rule conditions. The rule conditions determine the subset of objects to which the `Or` rule applies. Objects can match any of the listed filter conditions, which are joined by the `Or` logical operator. Only one of each filter condition is allowed.

Contents

MatchAnyPrefix

Filters objects that match any of the specified prefixes.

Type: Array of strings

Required: No

MatchAnySuffix

Filters objects that match any of the specified suffixes.

Type: Array of strings

Required: No

MatchAnyTag

Filters objects that match any of the specified S3 object tags.

Type: Array of [S3Tag](#) data types

Required: No

MatchObjectAge

Filters objects that match the specified object age range.

Type: [MatchObjectAge](#) data type

Required: No

MatchObjectSize

Filters objects that match the specified object size range.

Type: [MatchObjectSize](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

StorageLensTableDestination

Service: Amazon S3 Control

A container for configuring your S3 Storage Lens reports to export to read-only S3 table buckets. This parameter enables you to store your Storage Lens metrics in a structured, queryable table format in Apache Iceberg.

For more information about S3 Storage Lens, see [Assessing your storage activity and usage with S3 Storage Lens](#) in the *Amazon S3 User Guide*.

Contents

IsEnabled

A container that indicates whether the export to read-only S3 table buckets is enabled for your S3 Storage Lens configuration. When set to true, Storage Lens reports are automatically exported to tables in addition to other configured destinations.

Type: Boolean

Required: Yes

Encryption

A container for the encryption of the S3 Storage Lens metrics exports.

Type: [StorageLensDataExportEncryption](#) data type

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

StorageLensTag

Service: Amazon S3 Control

Contents

Key

Type: String

Length Constraints: Minimum length of 1. Maximum length of 128.

Pattern: `^([\p{L}\p{Z}\p{N}_.: /+=\ -@]*)$`

Required: Yes

Value

Type: String

Length Constraints: Minimum length of 0. Maximum length of 256.

Pattern: `^([\p{L}\p{Z}\p{N}_.: /+=\ -@]*)$`

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Tag

Service: Amazon S3 Control

A key-value pair that you use to label your resources. You can add tags to new resources when you create them, or you can add tags to existing resources. Tags can help you organize, track costs for, and control access to resources.

Contents

Key

The key of the key-value pair of a tag added to your AWS resource. A tag key can be up to 128 Unicode characters in length and is case-sensitive. System created tags that begin with `aws :` aren't supported.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 128.

Pattern: `^([\p{L}\p{Z}\p{N}_.: /+=\ -@]*)$`

Required: Yes

Value

The value of the key-value pair of a tag added to your AWS resource. A tag value can be up to 256 Unicode characters in length and is case-sensitive.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 256.

Pattern: `^([\p{L}\p{Z}\p{N}_.: /+=\ -@]*)$`

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Tagging

Service: Amazon S3 Control

Contents

TagSet

A collection for a set of tags.

Type: Array of [S3Tag](#) data types

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Transition

Service: Amazon S3 Control

Specifies when an object transitions to a specified storage class. For more information about Amazon S3 Lifecycle configuration rules, see [Transitioning objects using Amazon S3 Lifecycle](#) in the *Amazon S3 User Guide*.

Contents

Date

Indicates when objects are transitioned to the specified storage class. The date value must be in ISO 8601 format. The time is always midnight UTC.

Type: Timestamp

Required: No

Days

Indicates the number of days after creation when objects are transitioned to the specified storage class. The value must be a positive integer.

Type: Integer

Required: No

StorageClass

The storage class to which you want the object to transition.

Type: String

Valid Values: GLACIER | STANDARD_IA | ONEZONE_IA | INTELLIGENT_TIERING | DEEP_ARCHIVE

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

VersioningConfiguration

Service: Amazon S3 Control

Describes the versioning state of an Amazon S3 on Outposts bucket. For more information, see [PutBucketVersioning](#).

Contents

MFADelete

Specifies whether MFA delete is enabled or disabled in the bucket versioning configuration for the S3 on Outposts bucket.

Type: String

Valid Values: Enabled | Disabled

Required: No

Status

Sets the versioning state of the S3 on Outposts bucket.

Type: String

Valid Values: Enabled | Suspended

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

VpcConfiguration

Service: Amazon S3 Control

The virtual private cloud (VPC) configuration for an access point.

Contents

VpcId

If this field is specified, this access point will only allow connections from the specified VPC ID.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Amazon S3 Files

The following data types are supported by Amazon S3 Files:

- [CreationPermissions](#)
- [ExpirationDataRule](#)
- [ImportDataRule](#)
- [ListAccessPointsDescription](#)
- [ListFileSystemsDescription](#)
- [ListMountTargetsDescription](#)
- [PosixUser](#)

- [RootDirectory](#)
- [Tag](#)

CreationPermissions

Service: Amazon S3 Files

Specifies the permissions to set on newly created directories within the file system.

Contents

ownerGid

The POSIX group ID to assign to newly created directories.

Type: Long

Valid Range: Minimum value of 0. Maximum value of 4294967295.

Required: Yes

ownerUid

The POSIX user ID to assign to newly created directories.

Type: Long

Valid Range: Minimum value of 0. Maximum value of 4294967295.

Required: Yes

permissions

The octal permissions to assign to newly created directories.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 4.

Pattern: `[0-7]{3,4}`

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ExpirationDataRule

Service: Amazon S3 Files

Specifies a rule that controls when cached data expires from the file system based on last access time.

Contents

daysAfterLastAccess

The number of days after last access before cached data expires from the file system.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 365.

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ImportDataRule

Service: Amazon S3 Files

Specifies a rule that controls how data is imported from S3 into the file system.

Contents

prefix

The S3 key prefix that scopes this import rule. Only objects with keys beginning with this prefix are subject to the rule.

Type: String

Pattern: (| . * /)

Required: Yes

sizeLessThan

The upper size limit in bytes for this import rule. Only objects with a size strictly less than this value will have data imported into the file system.

Type: Long

Valid Range: Minimum value of 0. Maximum value of 52673613135872.

Required: Yes

trigger

The event that triggers data import. Valid values are ON_DIRECTORY_FIRST_ACCESS (import when a directory is first accessed) and ON_FILE_ACCESS (import when a file is accessed).

Type: String

Valid Values: ON_DIRECTORY_FIRST_ACCESS | ON_FILE_ACCESS

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ListAccessPointsDescription

Service: Amazon S3 Files

Contains information about an S3 File System Access Point returned in list operations.

Contents

accessPointArn

The Amazon Resource Name (ARN) of the access point.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 256.

Pattern: `arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}/access-point/fsap-[0-9a-f]{17,40}`

Required: Yes

accessPointId

The ID of the access point.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 256.

Pattern: `(arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}/access-point/fsap-[0-9a-f]{17,40}|fsap-[0-9a-f]{17,40})`

Required: Yes

fileSystemId

The ID of the S3 File System.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 128.

Pattern: `(arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}|fs-[0-9a-f]{17,40})`

Required: Yes

ownerId

The AWS account ID of the access point owner.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 12.

Pattern: $(\backslash d\{12\}) | (\backslash d\{4\} - \{4\} - \backslash d\{4\})$

Required: Yes

status

The current status of the access point.

Type: String

Valid Values: `available` | `creating` | `deleting` | `deleted` | `error` | `updating`

Required: Yes

name

The name of the access point.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 256.

Pattern: $([\backslash p\{L\}\backslash p\{Z\}\backslash p\{N\}_ . : / = + \backslash - @] ^ *)$

Required: No

posixUser

The POSIX identity configured for this access point.

Type: [PosixUser](#) object

Required: No

rootDirectory

The root directory configuration for this access point.

Type: [RootDirectory](#) object

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ListFileSystemsDescription

Service: Amazon S3 Files

Contains information about an S3 File System returned in list operations.

Contents

bucket

The Amazon Resource Name (ARN) of the S3 bucket.

Type: String

Pattern: (`arn:aws[a-zA-Z0-9-]*:s3:::.+`)

Required: Yes

creationTime

The time when the file system was created.

Type: Timestamp

Required: Yes

fileSystemArn

The Amazon Resource Name (ARN) of the file system.

Type: String

Pattern: (`arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}`)

Required: Yes

fileSystemId

The ID of the file system.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 128.

Pattern: (`arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}|fs-[0-9a-f]{17,40}`)

Required: Yes

ownerId

The AWS account ID of the file system owner.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 12.

Pattern: `(\d{12})|(\d{4}-\d{4}-\d{4})`

Required: Yes

roleArn

The Amazon Resource Name (ARN) of the IAM role used for S3 access.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 2048.

Pattern: `arn:(aws[a-zA-Z-]*)?:iam::\d{12}:role/?[a-zA-Z_0-9+=,.\@-_/]+`

Required: Yes

status

The current status of the file system.

Type: String

Valid Values: `available | creating | deleting | deleted | error | updating`

Required: Yes

name

The name of the file system.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 256.

Pattern: `([\p{L}\p{Z}\p{N}_.:/=+\-@]*)`

Required: No

statusMessage

Additional information about the file system status.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ListMountTargetsDescription

Service: Amazon S3 Files

Contains information about a mount target returned in list operations.

Contents

mountTargetId

The ID of the mount target.

Type: String

Length Constraints: Minimum length of 22. Maximum length of 45.

Pattern: fsmt-[0-9a-f]{17,40}

Required: Yes

ownerId

The AWS account ID of the mount target owner.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 12.

Pattern: (\d{12})|(\d{4}-\d{4}-\d{4})

Required: Yes

subnetId

The ID of the subnet where the mount target is located.

Type: String

Length Constraints: Minimum length of 15. Maximum length of 47.

Pattern: subnet-[0-9a-f]{8,40}

Required: Yes

availabilityZoneId

The Availability Zone ID where the mount target is located.

Type: String

Required: No

fileSystemId

The ID of the S3 File System.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 128.

Pattern: (arn:aws[-a-z]*:s3files:[0-9a-z-:]+:file-system/fs-[0-9a-f]{17,40}|fs-[0-9a-f]{17,40})

Required: No

ipv4Address

The IPv4 address of the mount target.

Type: String

Length Constraints: Minimum length of 7. Maximum length of 15.

Pattern: [0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}

Required: No

ipv6Address

The IPv6 address of the mount target.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 39.

Required: No

networkInterfaceId

The ID of the network interface associated with the mount target.

Type: String

Required: No

status

The current status of the mount target.

Type: String

Valid Values: available | creating | deleting | deleted | error | updating

Required: No

statusMessage

Additional information about the mount target status.

Type: String

Required: No

vpId

The ID of the VPC where the mount target is located.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

PosixUser

Service: Amazon S3 Files

Specifies the POSIX identity with uid, gid, and secondary group IDs for user enforcement.

Contents

gid

The POSIX group ID.

Type: Long

Valid Range: Minimum value of 0. Maximum value of 4294967295.

Required: Yes

uid

The POSIX user ID.

Type: Long

Valid Range: Minimum value of 0. Maximum value of 4294967295.

Required: Yes

secondaryGids

An array of secondary POSIX group IDs.

Type: Array of longs

Valid Range: Minimum value of 0. Maximum value of 4294967295.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Service: Amazon S3 Files

Contents

creationPermissions

Type: CreationPermissions object

path

Type: String

Pattern: (\\| (\\(?!\\.)+[^\\$#<>;`|&?{}^*\\/n]+){1,4})

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Tag

Service: Amazon S3 Files

A key-value pair for resource tagging.

Contents

key

The tag key. The key can't start with `aws :`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 128.

Pattern: `( [\p{L}\p{Z}\p{N}_ . : / = + \ - @ ] + )`

Required: Yes

value

The tag value.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 256.

Pattern: `( [\p{L}\p{Z}\p{N}_ . : / = + \ - @ ] * )`

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Amazon S3 on Outposts

The following data types are supported by Amazon S3 on Outposts:

- [Endpoint](#)
- [FailedReason](#)
- [NetworkInterface](#)
- [Outpost](#)

Endpoint

Service: Amazon S3 on Outposts

Amazon S3 on Outposts Access Points simplify managing data access at scale for shared datasets in S3 on Outposts. S3 on Outposts uses endpoints to connect to AWS Outposts buckets so that you can perform actions within your virtual private cloud (VPC). For more information, see [Accessing S3 on Outposts using VPC-only access points](#) in the *Amazon Simple Storage Service User Guide*.

Contents

AccessType

The type of connectivity used to access the Amazon S3 on Outposts endpoint.

Type: String

Valid Values: `Private` | `CustomerOwnedIp`

Required: No

CidrBlock

The VPC CIDR committed by this endpoint.

Type: String

Required: No

CreationTime

The time the endpoint was created.

Type: Timestamp

Required: No

CustomerOwnedIpv4Pool

The ID of the customer-owned IPv4 address pool used for the endpoint.

Type: String

Pattern: `^ipv4pool-coip-([0-9a-f]{17})$`

Required: No

EndpointArn

The Amazon Resource Name (ARN) of the endpoint.

Type: String

Pattern: `^arn:(aws|aws-cn|aws-us-gov|aws-iso|aws-iso-b):s3-outposts:[a-z\-0-9]*:[0-9]{12}:outpost/(op-[a-f0-9]{17}|ec2)/endpoint/[a-zA-Z0-9]{19}$`

Required: No

FailedReason

The failure reason, if any, for a create or delete endpoint operation.

Type: [FailedReason](#) object

Required: No

NetworkInterfaces

The network interface of the endpoint.

Type: Array of [NetworkInterface](#) objects

Required: No

OutpostsId

The ID of the AWS Outposts.

Type: String

Pattern: `^(op-[a-f0-9]{17}|\d{12}|ec2)$`

Required: No

SecurityGroupId

The ID of the security group used for the endpoint.

Type: String

Pattern: `^sg-([0-9a-f]{8}|[0-9a-f]{17})$`

Required: No

Status

The status of the endpoint.

Type: String

Valid Values: Pending | Available | Deleting | Create_Failed | Delete_Failed

Required: No

SubnetId

The ID of the subnet used for the endpoint.

Type: String

Pattern: ^subnet-([0-9a-f]{8}|[0-9a-f]{17})\$

Required: No

VpcId

The ID of the VPC used for the endpoint.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

FailedReason

Service: Amazon S3 on Outposts

The failure reason, if any, for a create or delete endpoint operation.

Contents

ErrorCode

The failure code, if any, for a create or delete endpoint operation.

Type: String

Required: No

Message

Additional error details describing the endpoint failure and recommended action.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

NetworkInterface

Service: Amazon S3 on Outposts

The container for the network interface.

Contents

NetworkInterfaceId

The ID for the network interface.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Outpost

Service: Amazon S3 on Outposts

Contains the details for the Outpost object.

Contents

CapacityInBytes

The Amazon S3 capacity of the outpost in bytes.

Type: Long

Required: No

OutpostArn

Specifies the unique Amazon Resource Name (ARN) for the outpost.

Type: String

Pattern: `^arn:(aws|aws-cn|aws-us-gov|aws-iso|aws-iso-b):outposts:[a-z\-\0-9]*:[0-9]{12}:outpost/(op-[a-f0-9]{17}|ec2)$`

Required: No

OutpostId

Specifies the unique identifier for the outpost.

Type: String

Pattern: `^(op-[a-f0-9]{17}|\d{12}|ec2)$`

Required: No

OwnerId

Returns the AWS account ID of the outpost owner. Useful for comparing owned versus shared outposts.

Type: String

Pattern: `^\d{12}$`

Required: No

S3OutpostArn

Specifies the unique S3 on Outposts ARN for use with AWS Resource Access Manager (AWS RAM).

Type: String

Pattern: `^arn:(aws|aws-cn|aws-us-gov|aws-iso|aws-iso-b):s3-outposts:[a-z\-\0-9]*:[0-9]{12}:outpost/(op-[a-f0-9]{17}|\d{12})/s3$`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Amazon S3 Tables

The following data types are supported by Amazon S3 Tables:

- [EncryptionConfiguration](#)
- [IcebergCompactionSettings](#)
- [IcebergMetadata](#)
- [IcebergPartitionField](#)
- [IcebergPartitionSpec](#)
- [IcebergSchema](#)
- [IcebergSchemaV2](#)
- [IcebergSnapshotManagementSettings](#)
- [IcebergSortField](#)
- [IcebergSortOrder](#)

- [IcebergUnreferencedFileRemovalSettings](#)
- [LastSuccessfulReplicatedUpdate](#)
- [ManagedTableInformation](#)
- [NamespaceSummary](#)
- [ReplicationDestination](#)
- [ReplicationDestinationStatusModel](#)
- [ReplicationInformation](#)
- [SchemaField](#)
- [SchemaV2Field](#)
- [StorageClassConfiguration](#)
- [TableBucketMaintenanceConfigurationValue](#)
- [TableBucketMaintenanceSettings](#)
- [TableBucketReplicationConfiguration](#)
- [TableBucketReplicationRule](#)
- [TableBucketSummary](#)
- [TableMaintenanceConfigurationValue](#)
- [TableMaintenanceJobStatusValue](#)
- [TableMaintenanceSettings](#)
- [TableMetadata](#)
- [TableRecordExpirationConfigurationValue](#)
- [TableRecordExpirationJobMetrics](#)
- [TableRecordExpirationSettings](#)
- [TableReplicationConfiguration](#)
- [TableReplicationRule](#)
- [TableSummary](#)

EncryptionConfiguration

Service: Amazon S3 Tables

Configuration specifying how data should be encrypted. This structure defines the encryption algorithm and optional KMS key to be used for server-side encryption.

Contents

sseAlgorithm

The server-side encryption algorithm to use. Valid values are AES256 for S3-managed encryption keys, or `aws:kms` for AWS KMS-managed encryption keys. If you choose SSE-KMS encryption you must grant the S3 Tables maintenance principal access to your KMS key. For more information, see [Permissions requirements for S3 Tables SSE-KMS encryption](#).

Type: String

Valid Values: AES256 | aws:kms

Required: Yes

kmsKeyArn

The Amazon Resource Name (ARN) of the KMS key to use for encryption. This field is required only when `sseAlgorithm` is set to `aws:kms`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: (`arn:aws[-a-z0-9]*:kms:[-a-z0-9]*:[0-9]{12}:key/.+`)

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)

- [AWS SDK for Ruby V3](#)

IcebergCompactionSettings

Service: Amazon S3 Tables

Contains details about the compaction settings for an Iceberg table.

Contents

strategy

The compaction strategy to use for the table. This determines how files are selected and combined during compaction operations.

Type: String

Valid Values: auto | binpack | sort | z-order

Required: No

targetFileSizeMB

The target file size for the table in MB.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 2147483647.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

IcebergMetadata

Service: Amazon S3 Tables

Contains details about the metadata for an Iceberg table.

Contents

partitionSpec

The partition specification for the Iceberg table. Partitioning organizes data into separate files based on the values of one or more fields, which can improve query performance by reducing the amount of data scanned. Each partition field applies a transform (such as identity, year, month, or bucket) to a single field.

Type: [IcebergPartitionSpec](#) object

Required: No

properties

A map of custom configuration properties for the Iceberg table.

Type: String to string map

Required: No

schema

The schema for an Iceberg table. Use this property to define table schemas with primitive types only. For schemas that include nested or complex types such as `struct`, `list`, or `map`, use `schemaV2` instead.

Type: [IcebergSchema](#) object

Required: No

schemaV2

The schema for an Iceberg table using the V2 format. Use this property to define table schemas that include nested or complex data types such as `struct`, `list`, or `map`, in addition to primitive types. For schemas with only primitive types, you can use either `schema` or `schemaV2`.

Type: [IcebergSchemaV2](#) object

Required: No

writeOrder

The sort order for the Iceberg table. Sort order defines how data is sorted within data files, which can improve query performance by enabling more efficient data skipping and filtering.

Type: [IcebergSortOrder](#) object

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

IcebergPartitionField

Service: Amazon S3 Tables

Defines a single partition field in an Iceberg partition specification.

Contents

name

The name for this partition field. This name is used in the partitioned file paths.

Type: String

Required: Yes

source-id

The ID of the source schema field to partition by. This must reference a valid field ID from the table schema.

Type: Integer

Required: Yes

transform

The partition transform to apply to the source field. Supported transforms include `identity`, `year`, `month`, `day`, `hour`, `bucket`, and `truncate`. For more information, see the [Apache Iceberg partition transforms documentation](#).

Type: String

Required: Yes

field-id

An optional unique identifier for this partition field. If not specified, S3 Tables automatically assigns a field ID.

Type: Integer

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

IcebergPartitionSpec

Service: Amazon S3 Tables

Defines how data in an Iceberg table is partitioned. Partitioning helps optimize query performance by organizing data into separate files based on field values. Each partition field specifies a transform to apply to a source field.

Contents

fields

The list of partition fields that define how the table data is partitioned. Each field specifies a source field and a transform to apply. This field is required if `partitionSpec` is provided.

Type: Array of [IcebergPartitionField](#) objects

Required: Yes

spec-id

The unique identifier for this partition specification. If not specified, defaults to 0.

Type: Integer

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

IcebergSchema

Service: Amazon S3 Tables

Contains details about the schema for an Iceberg table.

Contents

fields

The schema fields for the table

Type: Array of [SchemaField](#) objects

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

IcebergSchemaV2

Service: Amazon S3 Tables

Contains details about the schema for an Iceberg table using the V2 format. This schema format supports nested and complex data types such as `struct`, `list`, and `map`, in addition to primitive types.

Contents

fields

The schema fields for the table. Each field defines a column in the table, including its name, type, and whether it is required.

Type: Array of [SchemaV2Field](#) objects

Required: Yes

type

The type of the top-level schema, which is always a `struct` type as defined in the [Apache Iceberg specification](#). This value must be `struct`.

Type: String

Valid Values: `struct`

Required: Yes

identifier-field-ids

A list of field IDs that are used as the identifier fields for the table. Identifier fields uniquely identify a row in the table.

Type: Array of integers

Required: No

schema-id

An optional unique identifier for the schema. Schema IDs are used by Apache Iceberg to track schema evolution.

Type: Integer

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

IcebergSnapshotManagementSettings

Service: Amazon S3 Tables

Contains details about the snapshot management settings for an Iceberg table. The oldest snapshot expires when its age exceeds the `maxSnapshotAgeHours` and the total number of snapshots exceeds the value for the minimum number of snapshots to keep `minSnapshotsToKeep`.

Contents

maxSnapshotAgeHours

The maximum age of a snapshot before it can be expired.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 2147483647.

Required: No

minSnapshotsToKeep

The minimum number of snapshots to keep.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 2147483647.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

IcebergSortField

Service: Amazon S3 Tables

Defines a single sort field in an Iceberg sort order specification.

Contents

direction

The sort direction. Valid values are `asc` for ascending order or `desc` for descending order.

Type: String

Valid Values: `asc` | `desc`

Required: Yes

null-order

Specifies how null values are ordered. Valid values are `nulls-first` to place nulls before non-null values, or `nulls-last` to place nulls after non-null values.

Type: String

Valid Values: `nulls-first` | `nulls-last`

Required: Yes

source-id

The ID of the source schema field to sort by. This must reference a valid field ID from the table schema.

Type: Integer

Required: Yes

transform

The transform to apply to the source field before sorting. Use `identity` to sort by the field value directly, or specify other transforms as needed.

Type: String

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

IcebergSortOrder

Service: Amazon S3 Tables

Defines the sort order for data within an Iceberg table. Sorting data can improve query performance by enabling more efficient data skipping.

Contents

fields

The list of sort fields that define how data is sorted within files. Each field specifies a source field, sort direction, and null ordering. This field is required if `writeOrder` is provided.

Type: Array of [IcebergSortField](#) objects

Required: Yes

order-id

The unique identifier for this sort order. If not specified, defaults to 1. The order ID is used by Apache Iceberg to track sort order evolution.

Type: Integer

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

IcebergUnreferencedFileRemovalSettings

Service: Amazon S3 Tables

Contains details about the unreferenced file removal settings for an Iceberg table bucket.

Contents

nonCurrentDays

The number of days an object has to be non-current before it is deleted.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 2147483647.

Required: No

unreferencedDays

The number of days an object has to be unreferenced before it is marked as non-current.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 2147483647.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

LastSuccessfulReplicatedUpdate

Service: Amazon S3 Tables

Contains information about the most recent successful replication update to a destination.

Contents

metadataLocation

The S3 location of the metadata that was successfully replicated.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: Yes

timestamp

The timestamp when the replication update completed successfully.

Type: Timestamp

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ManagedTableInformation

Service: Amazon S3 Tables

Contains information about tables that are managed by S3 Tables, including replication information for replica tables.

Contents

replicationInformation

If this table is a replica, contains information about the source table from which it is replicated.

Type: [ReplicationInformation](#) object

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

NamespaceSummary

Service: Amazon S3 Tables

Contains details about a namespace.

Contents

createdAt

The date and time the namespace was created at.

Type: Timestamp

Required: Yes

createdBy

The ID of the account that created the namespace.

Type: String

Length Constraints: Fixed length of 12.

Pattern: `[0-9].*`

Required: Yes

namespace

The name of the namespace.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[0-9a-z_]*`

Required: Yes

ownerAccountId

The ID of the account that owns the namespace.

Type: String

Length Constraints: Fixed length of 12.

Pattern: `[0-9].*`

Required: Yes

namespaceId

The system-assigned unique identifier for the namespace.

Type: String

Required: No

tableBucketId

The system-assigned unique identifier for the table bucket that contains this namespace.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ReplicationDestination

Service: Amazon S3 Tables

Specifies a destination table bucket for replication.

Contents

destinationTableBucketARN

The Amazon Resource Name (ARN) of the destination table bucket where tables will be replicated.

Type: String

Pattern: (arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ReplicationDestinationStatusModel

Service: Amazon S3 Tables

Contains status information for a replication destination, including the current replication state, last successful update, and any error messages.

Contents

destinationTableBucketArn

The Amazon Resource Name (ARN) of the destination table bucket.

Type: String

Pattern: (`arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63}`)

Required: Yes

replicationStatus

The current status of replication to this destination.

Type: String

Valid Values: `pending` | `completed` | `failed`

Required: Yes

destinationTableArn

The Amazon Resource Name (ARN) of the destination table.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: (`arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63}/table/[a-zA-Z0-9_-]{1,255}`)

Required: No

failureMessage

If replication has failed, this field contains an error message describing the failure reason.

Type: String

Required: No

lastSuccessfulReplicatedUpdate

Information about the most recent successful replication update to this destination.

Type: [LastSuccessfulReplicatedUpdate](#) object

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ReplicationInformation

Service: Amazon S3 Tables

Contains information about the source of a replicated table.

Contents

sourceTableARN

The Amazon Resource Name (ARN) of the source table from which this table is replicated.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: (arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63}/table/[a-zA-Z0-9_-]{1,255})

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SchemaField

Service: Amazon S3 Tables

Contains details about a schema field.

Contents

name

The name of the field.

Type: String

Required: Yes

type

The field type. S3 Tables supports all Apache Iceberg primitive types. For more information, see the [Apache Iceberg documentation](#).

Type: String

Required: Yes

id

An optional unique identifier for the schema field. Field IDs are used by Apache Iceberg to track schema evolution and maintain compatibility across schema changes. If not specified, S3 Tables automatically assigns field IDs.

Type: Integer

Required: No

required

A Boolean value that specifies whether values are required for each row in this field. By default, this is `false` and null values are allowed in the field. If this is `true` the field does not allow null values.

Type: Boolean

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SchemaV2Field

Service: Amazon S3 Tables

Contains details about a schema field in the V2 format. This field format supports nested and complex data types such as `struct`, `list`, and `map`, in addition to primitive types.

Contents

id

The unique identifier for the schema field. Field IDs are used by Apache Iceberg to track schema evolution and maintain compatibility across schema changes.

Type: Integer

Required: Yes

name

The name of the field.

Type: String

Required: Yes

required

A Boolean value that specifies whether values are required for each row in this field. If this is `true`, the field does not allow null values.

Type: Boolean

Required: Yes

type

The data type of the field. This can be a primitive type string such as `boolean`, `int`, `long`, `float`, `double`, `string`, `binary`, `date`, `timestamp`, or `timestampz`, or a complex type represented as a JSON object for nested types such as `struct`, `list`, or `map`. For more information, see the [Apache Iceberg schemas and data types documentation](#).

Type: JSON value

Required: Yes

doc

An optional description of the field.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

StorageClassConfiguration

Service: Amazon S3 Tables

The configuration details for the storage class of tables or table buckets. This allows you to optimize storage costs by selecting the appropriate storage class based on your access patterns and performance requirements.

Contents

storageClass

The storage class for the table or table bucket. Valid values include storage classes optimized for different access patterns and cost profiles.

Type: String

Valid Values: STANDARD | INTELLIGENT_TIERING

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

TableBucketMaintenanceConfigurationValue

Service: Amazon S3 Tables

Details about the values that define the maintenance configuration for a table bucket.

Contents

settings

Contains details about the settings of the maintenance configuration.

Type: [TableBucketMaintenanceSettings](#) object

Note: This object is a Union. Only one member of this object can be specified or returned.

Required: No

status

The status of the maintenance configuration.

Type: String

Valid Values: enabled | disabled

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

TableBucketMaintenanceSettings

Service: Amazon S3 Tables

Contains details about the maintenance settings for the table bucket.

Contents

Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

icebergUnreferencedFileRemoval

The unreferenced file removal settings for the table bucket.

Type: [IcebergUnreferencedFileRemovalSettings](#) object

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

TableBucketReplicationConfiguration

Service: Amazon S3 Tables

The replication configuration for a table bucket. This configuration defines how tables in the source bucket are replicated to destination table buckets, including the IAM role used for replication.

Contents

role

The Amazon Resource Name (ARN) of the IAM role that S3 Tables assumes to replicate tables on your behalf.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `arn:.*:iam::[0-9]{12}:role/.+`

Required: Yes

rules

An array of replication rules that define which tables to replicate and where to replicate them.

Type: Array of [TableBucketReplicationRule](#) objects

Array Members: Fixed number of 1 item.

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

TableBucketReplicationRule

Service: Amazon S3 Tables

Defines a rule for replicating tables from a source table bucket to one or more destination table buckets.

Contents

destinations

An array of destination table buckets where tables should be replicated.

Type: Array of [ReplicationDestination](#) objects

Array Members: Minimum number of 1 item. Maximum number of 5 items.

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

TableBucketSummary

Service: Amazon S3 Tables

Contains details about a table bucket.

Contents

arn

The Amazon Resource Name (ARN) of the table bucket.

Type: String

Pattern: (arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63})

Required: Yes

createdAt

The date and time the table bucket was created at.

Type: Timestamp

Required: Yes

name

The name of the table bucket.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Pattern: [0-9a-z-]*

Required: Yes

ownerAccountId

The ID of the account that owns the table bucket.

Type: String

Length Constraints: Fixed length of 12.

Pattern: `[0-9].*`

Required: Yes

tableBucketId

The system-assigned unique identifier for the table bucket.

Type: String

Required: No

type

The type of the table bucket.

Type: String

Valid Values: `customer` | `aws`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

TableMaintenanceConfigurationValue

Service: Amazon S3 Tables

The values that define a maintenance configuration for a table.

Contents

settings

Contains details about the settings for the maintenance configuration.

Type: [TableMaintenanceSettings](#) object

Note: This object is a Union. Only one member of this object can be specified or returned.

Required: No

status

The status of the maintenance configuration.

Type: String

Valid Values: enabled | disabled

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

TableMaintenanceJobStatusValue

Service: Amazon S3 Tables

Details about the status of a maintenance job.

Contents

status

The status of the job.

Type: String

Valid Values: Not_Yet_Run | Successful | Failed | Disabled

Required: Yes

failureMessage

The failure message of a failed job.

Type: String

Required: No

lastRunTimestamp

The date and time that the maintenance job was last run.

Type: Timestamp

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

TableMaintenanceSettings

Service: Amazon S3 Tables

Contains details about maintenance settings for the table.

Contents

Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

icebergCompaction

Contains details about the Iceberg compaction settings for the table.

Type: [IcebergCompactionSettings](#) object

Required: No

icebergSnapshotManagement

Contains details about the Iceberg snapshot management settings for the table.

Type: [IcebergSnapshotManagementSettings](#) object

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

TableMetadata

Service: Amazon S3 Tables

Contains details about the table metadata.

Contents

Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

iceberg

Contains details about the metadata of an Iceberg table.

Type: [IcebergMetadata](#) object

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

TableRecordExpirationConfigurationValue

Service: Amazon S3 Tables

The expiration configuration settings for records in a table, and the status of the configuration. If the status of the configuration is enabled, records expire and are automatically removed after the number of days specified in the record expiration settings for the table.

Contents

settings

The expiration settings for records in the table.

Type: [TableRecordExpirationSettings](#) object

Required: No

status

The status of the expiration settings for records in the table.

Type: String

Valid Values: enabled | disabled

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

TableRecordExpirationJobMetrics

Service: Amazon S3 Tables

Provides metrics for the record expiration job that most recently ran for a table. The metrics provide insight into the amount of data that was removed when the job ran.

Contents

deletedDataFiles

The total number of data files that were removed when the job ran.

Type: Long

Required: No

deletedRecords

The total number of records that were removed when the job ran.

Type: Long

Required: No

removedFileSize

The total size (in bytes) of the data files that were removed when the job ran.

Type: Long

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

TableRecordExpirationSettings

Service: Amazon S3 Tables

The record expiration setting that specifies when records expire and are automatically removed from a table.

Contents

days

If you enable record expiration for a table, you can specify the number of days to retain your table records. For example, to retain your table records for one year, set this value to 365.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 2147483647.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

TableReplicationConfiguration

Service: Amazon S3 Tables

The replication configuration for an individual table. This configuration defines how the table is replicated to destination tables.

Contents

role

The Amazon Resource Name (ARN) of the IAM role that S3 Tables assumes to replicate the table on your behalf.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `arn:.*:iam::[0-9]{12}:role/.+`

Required: Yes

rules

An array of replication rules that define where this table should be replicated.

Type: Array of [TableReplicationRule](#) objects

Array Members: Fixed number of 1 item.

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

TableReplicationRule

Service: Amazon S3 Tables

Defines a rule for replicating a table to one or more destination tables.

Contents

destinations

An array of destination table buckets where this table should be replicated.

Type: Array of [ReplicationDestination](#) objects

Array Members: Minimum number of 1 item. Maximum number of 5 items.

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

TableSummary

Service: Amazon S3 Tables

Contains details about a table.

Contents

createdAt

The date and time the table was created at.

Type: Timestamp

Required: Yes

modifiedAt

The date and time the table was last modified at.

Type: Timestamp

Required: Yes

name

The name of the table.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: [0-9a-z_]*

Required: Yes

namespace

The name of the namespace.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: [0-9a-z_]*

Required: Yes

tableARN

The Amazon Resource Name (ARN) of the table.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: (`arn:aws[-a-z0-9]*:[a-z0-9]+:[-a-z0-9]*:[0-9]{12}:bucket/[a-z0-9_-]{3,63}/table/[a-zA-Z0-9_-]{1,255}`)

Required: Yes

type

The type of the table.

Type: String

Valid Values: `customer` | `aws`

Required: Yes

managedByService

The AWS service managing this table, if applicable. For example, a replicated table is managed by the S3 Tables replication service.

Type: String

Required: No

namespaceId

The unique identifier for the namespace that contains this table.

Type: String

Required: No

tableBucketId

The unique identifier for the table bucket that contains this table.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Amazon S3 Vectors

The following data types are supported by Amazon S3 Vectors:

- [EncryptionConfiguration](#)
- [GetOutputVector](#)
- [Index](#)
- [IndexSummary](#)
- [ListOutputVector](#)
- [MetadataConfiguration](#)
- [PutInputVector](#)
- [QueryOutputVector](#)
- [ValidationExceptionField](#)
- [VectorBucket](#)
- [VectorBucketSummary](#)
- [VectorData](#)

EncryptionConfiguration

Service: Amazon S3 Vectors

The encryption configuration for a vector bucket or index. By default, if you don't specify, all new vectors in Amazon S3 vector buckets use server-side encryption with Amazon S3 managed keys (SSE-S3), specifically AES256. You can optionally override bucket level encryption settings, and set a specific encryption configuration for a vector index at the time of index creation.

Contents

kmsKeyArn

AWS Key Management Service (KMS) customer managed key ID to use for the encryption configuration. This parameter is allowed if and only if `sseType` is set to `aws:kms`.

To specify the KMS key, you must use the format of the KMS key Amazon Resource Name (ARN).

For example, specify Key ARN in the following format: `arn:aws:kms:us-east-2:111122223333:key/1234abcd-12ab-34cd-56ef-1234567890ab`

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: `(arn:aws[-a-z0-9]*:kms:[-a-z0-9]*:[0-9]{12}:key/.+)`

Required: No

sseType

The server-side encryption type to use for the encryption configuration of the vector bucket. By default, if you don't specify, all new vectors in Amazon S3 vector buckets use server-side encryption with Amazon S3 managed keys (SSE-S3), specifically AES256.

Type: String

Valid Values: `AES256` | `aws:kms`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

GetOutputVector

Service: Amazon S3 Vectors

The attributes of a vector returned by the GetVectors operation.

Contents

key

The name of the vector.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

data

The vector data of the vector.

Type: [VectorData](#) object

Note: This object is a Union. Only one member of this object can be specified or returned.

Required: No

metadata

Metadata about the vector.

Type: JSON value

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Index

Service: Amazon S3 Vectors

The attributes of a vector index.

Contents

creationTime

Date and time when the vector index was created.

Type: Timestamp

Required: Yes

dataType

The data type of the vectors inserted into the vector index.

Type: String

Valid Values: float32

Required: Yes

dimension

The number of values in the vectors that are inserted into the vector index.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 4096.

Required: Yes

distanceMetric

The distance metric to be used for similarity search.

Type: String

Valid Values: euclidean | cosine

Required: Yes

indexArn

The Amazon Resource Name (ARN) of the vector index.

Type: String

Pattern: `arn:aws[-a-z0-9]*:s3vectors:[a-z0-9-]+:[0-9]{12}:bucket/[a-z0-9][a-z0-9- .]{1,61}[a-z0-9]/index/[a-z0-9][a-z0-9- .]{1,61}[a-z0-9]`

Required: Yes

indexName

The name of the vector index.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: Yes

vectorBucketName

The name of the vector bucket that contains the vector index.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: Yes

encryptionConfiguration

The encryption configuration for a vector index. By default, if you don't specify, all new vectors in the vector index will use the encryption configuration of the vector bucket.

Type: [EncryptionConfiguration](#) object

Required: No

metadataConfiguration

The metadata configuration for the vector index.

Type: [MetadataConfiguration](#) object

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

IndexSummary

Service: Amazon S3 Vectors

Summary information about a vector index.

Contents

creationTime

Date and time when the vector index was created.

Type: Timestamp

Required: Yes

indexArn

The Amazon Resource Name (ARN) of the vector index.

Type: String

Pattern: `arn:aws[-a-z0-9]*:s3vectors:[a-z0-9-]+:[0-9]{12}:bucket/[a-z0-9][a-z0-9-]{1,61}[a-z0-9]/index/[a-z0-9][a-z0-9-]{1,61}[a-z0-9]`

Required: Yes

indexName

The name of the vector index.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: Yes

vectorBucketName

The name of the vector bucket that contains the vector index.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ListOutputVector

Service: Amazon S3 Vectors

The attributes of a vector returned by the `ListVectors` operation.

Contents

key

The name of the vector.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

data

The vector data of the vector.

Type: [VectorData](#) object

Note: This object is a Union. Only one member of this object can be specified or returned.

Required: No

metadata

Metadata about the vector.

Type: JSON value

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

MetadataConfiguration

Service: Amazon S3 Vectors

The metadata configuration for a vector index.

Contents

nonFilterableMetadataKeys

Non-filterable metadata keys allow you to enrich vectors with additional context during storage and retrieval. Unlike default metadata keys, these keys can't be used as query filters. Non-filterable metadata keys can be retrieved but can't be searched, queried, or filtered. You can access non-filterable metadata keys of your vectors after finding the vectors. For more information about non-filterable metadata keys, see [Vectors](#) and [Limitations and restrictions](#) in the *Amazon S3 User Guide*.

Type: Array of strings

Array Members: Minimum number of 1 item. Maximum number of 10 items.

Length Constraints: Minimum length of 1. Maximum length of 63.

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

PutInputVector

Service: Amazon S3 Vectors

The attributes of a vector to add to a vector index.

Contents

data

The vector data of the vector.

Vector dimensions must match the dimension count that's configured for the vector index.

- For the `cosine` distance metric, zero vectors (vectors containing all zeros) aren't allowed.
- For both `cosine` and `euclidean` distance metrics, vector data must contain only valid floating-point values. Invalid values such as NaN (Not a Number) or Infinity aren't allowed.

Type: [VectorData](#) object

Note: This object is a Union. Only one member of this object can be specified or returned.

Required: Yes

key

The name of the vector. The key uniquely identifies the vector in a vector index.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

metadata

Metadata about the vector. All metadata entries undergo validation to ensure they meet the format requirements for size and data types.

Type: JSON value

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

QueryOutputVector

Service: Amazon S3 Vectors

The attributes of a vector in the approximate nearest neighbor search.

Contents

key

The key of the vector in the approximate nearest neighbor search.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Required: Yes

distance

The measure of similarity between the vector in the response and the query vector.

Type: Float

Required: No

metadata

The metadata associated with the vector, if requested.

Type: JSON value

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ValidationExceptionField

Service: Amazon S3 Vectors

Contains information about a validation exception.

Contents

message

A message about the validation exception.

Type: String

Required: Yes

path

A path about the validation exception.

Type: String

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

VectorBucket

Service: Amazon S3 Vectors

The attributes of a vector bucket.

Contents

creationTime

Date and time when the vector bucket was created.

Type: Timestamp

Required: Yes

vectorBucketArn

The Amazon Resource Name (ARN) of the vector bucket.

Type: String

Pattern: `arn:aws[-a-z0-9]*:s3vectors:[a-z0-9-]+:[0-9]{12}:bucket/[a-z0-9][a-z0-9-.]{1,61}[a-z0-9]`

Required: Yes

vectorBucketName

The name of the vector bucket.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: Yes

encryptionConfiguration

The encryption configuration for the vector bucket.

Type: [EncryptionConfiguration](#) object

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

VectorBucketSummary

Service: Amazon S3 Vectors

Summary information about a vector bucket.

Contents

creationTime

Date and time when the vector bucket was created.

Type: Timestamp

Required: Yes

vectorBucketArn

The Amazon Resource Name (ARN) of the vector bucket.

Type: String

Pattern: `arn:aws[-a-z0-9]*:s3vectors:[a-z0-9-]+:[0-9]{12}:bucket/[a-z0-9][a-z0-9-.]{1,61}[a-z0-9]`

Required: Yes

vectorBucketName

The name of the vector bucket.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)

- [AWS SDK for Ruby V3](#)

VectorData

Service: Amazon S3 Vectors

The vector data in different formats.

Contents

Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

float32

The vector data as 32-bit floating point numbers. The number of elements in this array must exactly match the dimension of the vector index where the operation is being performed.

Type: Array of floats

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)