



User Guide

AWS Transform MGN



AWS Transform MGN: User Guide

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

What Is AWS Transform MGN?	1
Accessing the AWS Transform MGN console	1
Supported AWS Regions	1
MGN technical training materials	4
Quick start guide	5
First time setup	5
Adding source servers	6
Migration lifecycle	7
Configuring launch settings	8
Launching a test instance	9
Ready for testing indicators	10
Starting a test	10
Reverting or finalizing a test	11
Launching a cutover instance	13
Ready for cutover indicators	13
Starting a cutover	14
Reverting or finalizing a cutover	14
Getting started	17
Migration workflow	17
Initializing with the console	18
IAM role creation	19
Additional policies	19
Initializing with the API	20
Creating the required IAM roles	21
Creating the templates	23
Configuring the templates	23
Configuring your replication template	23
Configuring your launch template	23
Configuring a post-launch template	24
Using the console	24
Source servers page	24
Best practices	26
Planning	26
Testing	26

Successful implementation	26
Ensuring project success	27
Migration at scale	27
Choose your agent-based vs agentless migration approach	27
Cost planning	27
Import/Export	28
MGN Connector	28
Network benchmarking	28
Storage benchmarking	29
Dedicated replication servers	29
Wave planning	30
Cutover planning	30
Clustered servers	30
Environment updates/changes	31
API limits	31
Service quota limits	32
Account settings	37
Replication template	37
Understanding template settings and server-specific settings	38
Edit your replication settings template	38
Edit replication settings for a server	38
Replication template reference	39
Launch template	47
General launch settings reference	47
Default EC2 launch template settings	49
MAP program tagging setting	50
Post-launch template	50
Activating post-launch settings	51
Editing the post-launch settings template	52
Deploying post-launch actions	53
Encrypt post-launch action parameters	53
Post-launch actions table	53
Create a custom post-launch action	53
Edit custom post-launch actions	55
Predefined post-launch actions reference	55
Network requirements	65

Service and network architecture overview	65
Network setting preparations	66
Staging area subnet	66
Operational subnets	67
Active Directory connectivity	67
Required connectivity settings	68
Endpoints and firewall allowlist	69
Communication between source servers and AWS Transform MGN over TCP port 443	71
Communication between source servers and the staging area subnet over TCP port 1500	71
Communication between the staging area subnet and AWS Transform MGN over TCP port 443	72
Communication between the staging area subnet and Amazon S3	72
Network requirements for FSx for ONTAP	73
Troubleshooting connectivity issues	75
Enabling Windows Firewall for TCP port 443 connectivity	75
Enabling Linux Firewall for TCP port 443 connectivity	77
Source servers	79
Adding source servers	80
Installation requirements	80
Supported operating systems	87
Installing the AWS Replication Agent	104
Installing the vCenter Client	262
Agentless replication overview	263
VMware limitations	265
Generating vCenter Client IAM credentials	265
Installing the MGN vCenter Client	266
Replicating servers from vCenter to AWS	272
Updating the vCenter or AWS Credentials	273
Differentiating agentless and agent-based servers	274
Manage source servers	275
Interacting with the source servers page	275
Add a source server	277
Manage data replication	280
Manage test and cutover instances	280
Source server migration metrics	283

Filtering the source servers list	284
Server details	285
Migration dashboard	286
Server info	298
Tags	299
Disk settings	300
Change target storage type	300
Replication settings	301
Launch settings	301
Post-launch settings	301
Editing the post-launch settings	302
Activating and deactivating post-launch actions	303
Deploying post-launch actions	304
Source server actions	305
Requirements	305
Actions before migration	305
Edit replication settings	306
Edit launch settings	306
Add or manage tags	307
Configure post-launch actions	307
Assign server to an application	307
Actions during migration	308
Start data replication	308
Pause data replication	308
Resume data replication	309
Launch test instance	309
Finalize test and mark as ready for cutover	310
Launch cutover instance	310
Monitor replication health	311
Actions after migration	311
Finalize cutover	311
Disconnect from AWS Transform MGN	312
Archive source server	312
Delete source server	313
Source server actions summary	314
Applications	316

Applications page	316
Interacting with the Applications page	316
Add application	317
Edit application	318
Delete application	318
Manage applications	318
Filtering	320
Application details	321
Overview dashboard	323
Review aggregated source server migration metrics	324
Source servers table	328
Tags	328
Waves	329
Waves page	329
Interacting with the Waves page	329
Add wave	330
Edit wave	331
Delete wave	331
Actions menu	331
Filtering	333
Wave details	333
Overview dashboard	335
Applications	336
Source servers	338
MGN Connectors	344
Prerequisites	344
General prerequisites	344
Operating systems that support the MGN connector	345
SSM requirements	345
Security recommendations	345
Architecture overview	345
Setup instructions	346
Connector reuse and lifecycle	348
IAM roles for connector	348
Create roles using the console	349
Create roles manually	356

Deploy role using CloudFormation template	360
Installing the MGN connector on a secured network	360
Global view	361
Create VPC endpoints	361
Create a Route 53 inbound endpoint	362
Modify replication settings	363
Verify VPC endpoints are being used	363
Review details about your MGN connectors	363
MGN Connectors management	364
Introduction to the MGN connector page	364
Edit connector	365
Delete MGN connector	365
Register server credentials	366
Verify source server prerequisites	368
Install the replication agent	368
View command history	369
Import and export	370
Configuration editing	370
Deleting values from the inventory file	371
Deleting tags from the inventory file	371
Importing your data inventory	371
Required permissions	372
Inventory Import parameters	373
Import from a local disk	381
Importing from an S3 bucket	382
Import history	382
Exporting your data inventory	382
Required permissions	383
Amazon S3 bucket permission requirements	384
Export parameters	384
Export to a local disk	385
Exporting to an S3 bucket	385
Export history	386
Global view	387
Setting up your AWS Organization	387
Activate trusted access	387

Setting up StackSets	388
Using an AWS KMS customer managed key for encryption	389
Inviting other accounts	390
Using global view	390
Source servers in member accounts	391
Applications	392
Waves	392
Import/Export	393
Launching test and cutover instances	394
Preparing for test and cutover instance launch	394
Launch settings	395
General launch settings	396
EC2 launch template	401
Launching test instances	410
Ready for testing indicators	411
Starting a test	411
Reverting a test	412
Marking as Ready for cutover	413
Launching cutover instances	414
Ready for cutover indicators	414
Starting a cutover	415
Reverting a cutover	416
Finalizing a cutover	416
Launch history	418
Overview	418
Job details	418
Target storage	420
Amazon EBS storage	420
Replication template settings	420
Server-level settings	425
FSx for ONTAP	426
Overview	426
Prerequisites	428
Step 1: Configure security groups	428
Step 2: Create FSx for ONTAP file system	431
Step 3: Configure certificate-based authentication	433

Step 4: Store certificates in AWS Secrets Manager	436
Step 5: Configure MGN replication settings	437
Step 6: Configure launch template and launch settings	438
Step 7: Enable volume integrity validation (recommended)	438
Post-migration optimization	438
Generating a self-signed certificate	438
Monitoring AWS Transform MGN	443
Monitoring with CloudWatch	443
MGN and EventBridge	444
Event samples	445
Registering event rules	447
Logging AWS Transform MGN with AWS CloudTrail	448
AWS Transform MGN information in CloudTrail	449
Understanding AWS Transform MGN log file entries	450
Security	452
Overview	452
Identity and access management	453
Federated identity	454
Authenticating with identities	454
Grant permission to tag resources during creation	455
AWS managed policies	456
Managing access using policies	474
Identity-based policies	474
Using identity-based policies	475
Resource-based policies	481
Access control lists (ACLs)	481
Other policy types	482
Multiple policy types	482
Using service-linked roles	482
AWSServiceRoleForApplicationMigrationService service-linked role	483
Creating a service-linked role for AWS Transform MGN	483
Editing a service-linked role for AWS Transform MGN	484
Deleting a service-linked role for AWS Transform MGN	484
Supported Regions for MGN service-linked roles	486
Policy structure	486
Resilience	487

Infrastructure security	487
Compliance validation	489
Cross-service confused deputy prevention	489
Troubleshooting	492
Troubleshooting launch errors	492
Do I need to recreate the entire launch template for every version?	492
Error - AccessDeniedException - Must be admin user	493
VPCIdNotSpecified error	493
Error: Failed to connect using HTTP channel	493
Could not take up-to-date snapshot. Launching from snapshot taken on... ..	494
Troubleshooting communication errors	494
Solving communication problems over TCP Port 443 between the staging area and the AWS Transform MGN	495
Authenticate with service errors	496
Calculating the required bandwidth for TCP Port 1500	496
Verifying communication over Port 1500	497
Solving communication problems over Port 1500	498
Missing Amazon Linux 2023 repository access	500
Troubleshooting agent issues	501
Error: Installation failed	501
Where can I find the MGN Agent logs?	509
Troubleshooting agentless replication issues	509
Common replication errors	511
Agent not seen	511
Not converging	512
Snapshot failure	512
Unstable network	512
Failed to connect AWS replication Agent to replication software	513
Failed to establish communication with replication software	513
Failed to create firewall rules	513
Failed to create security groups	513
Failed to authenticate with service	513
Failed to create staging disks	514
Replication stuck at Attach Staging Disks	514
Failed to pair the replication agent with replication server	515
Stalled replication when replicating a source volume smaller than 1MiB	515

Unknown data replication error	515
Other troubleshooting topics	515
Re-initialize the AWS Transform MGN	516
Windows license activation – AWS	516
Migration leaving behind replication volumes after cutover	516
Replication lag issues	517
Windows drive letter changes	517
Windows Dynamic Disk troubleshooting	517
Deleting MGN resources	518
Set UEFI boot mode	519
Target instance cannot connect to Active Directory after migration	519
Troubleshooting FSx for ONTAP issues	520
Troubleshooting FSx for ONTAP iSCSI connectivity	520
FSx for ONTAP replication errors	522
Replication volume not deleted after Finalize cutover/Disconnect from service (FlexClone split blocked by backup)	524
Orphaned FSx for ONTAP target volumes (FlexClone) after launch cleanup	525
Troubleshooting FSx for ONTAP launch errors	526
FAQ	529
General questions	529
Why was AWS Application Migration Service renamed to AWS Transform MGN?	530
Can MGN protect or migrate physical servers?	530
What data is stored on and transmitted through MGN service?	530
What should I consider when replicating Active Directory?	530
Does AWS Transform MGN work with LVM and RAID configurations?	531
What is there to note regarding SAN/NAS support?	531
Does AWS Transform MGN support Windows License migration?	531
Can you perform an OS (Operating System) upgrade with AWS Transform MGN?	531
What are the AWS Transform MGN quota limits?	531
What are the Private APIs used by MGN to define actions in the IAM Policy?	534
Which post-launch scripts does MGN support?	535
What happens if I use a custom DNS?	536
Can I use AWS Transform MGN to migrate servers from VMware Cloud on AWS (VMC) to Amazon EC2?	536
When should I use AWS Elastic Disaster Recovery (AWS DRS) for migration?	537
Agent related	537

What does the AWS Replication Agent do?	538
What kind of data is transferred between the agent and the AWS Transform MGN?	538
Can a proxy server be used between the source server and the AWS Transform MGN console?	539
What are the prerequisites needed to install the AWS Replication Agent?	539
What ports does the AWS Replication Agent use?	540
What privileges does the AWS Replication Agent require?	540
Is it possible to install the agent on servers running operating systems that are not listed as supported?	540
What kind of resources does the AWS Replication Agent use?	540
Can AWS Transform MGN migrate containers?	540
Does the AWS Replication Agent cache any data to disk?	540
How is communication between the AWS Replication Agent and the AWS Transform MGN secured?	541
Is it possible to change the port the AWS Replication Agent uses from TCP Port 1500 to a different port?	541
How do I manually uninstall the AWS Transform MGN agent from a server?	541
When do I need to reinstall the agent?	541
How much bandwidth does the AWS Replication Agent consume?	542
How many disks can the AWS Replication Agent replicate?	542
Is it possible to add a disk to replication without a complete resync of any disks that have already been replicated?	542
Is the AWS Replication Agent installed on launched test and cutover instances?	542
How do temporary credentials work?	543
Which Windows and Linux OSs support no-rescan upon reboot?	543
Agentless replication related	544
In which situations would you recommend using agentless replication (snapshot shipping)?	545
In which situations would you recommend using agent-based replication?	545
How does agentless replication work?	545
Does agentless replication require installing any component in the customer's source data center?	545
Is the agentless feature available in all Regions that MGN service supports?	545
Does agentless replication support the same source operating systems that are supported by agent-based replication?	546
Is the agentless feature supported in CloudEndure migration?	546

Which virtualization environments are supported by the agentless feature?	546
On which operating systems can the MGN vCenter Client be installed?	546
Do I need to generate special credentials to install the MGN vCenter Client?	546
What are the agentless replication prerequisites?	546
How do I install the MGN vCenter Client?	547
Can a proxy server be used between the source server and the AWS Transform MGN console?	547
Replication related	547
What is the lifecycle of the snapshots and volumes automatically created during migration?	548
What do Lag and Backlog mean during replication?	548
What is Continuous, Block level data replication?	549
What are the Replication initiation Steps?	549
Is the replicated data encrypted?	550
How is the Replication Server provisioned and managed in the Staging Area?	550
What type of replication server is utilized in the AWS Transform MGN staging area?	550
Can we set specific IP addresses for the replication server or conversion server in the AWS Transform MGN staging area?	550
Does AWS Transform MGN compress data during replication?	551
Are events that are generated by the AWS Transform MGN servers logged in Cloudtrail in AWS?	551
How many snapshots does AWS Transform MGN create?	551
Does AWS Transform MGN delete snapshots?	551
How much capacity is allocated to the staging area?	551
Why is 0.0.0.0:1500 added to inbound rules in the staging area?	551
Can AWS Transform MGN replicate Oracle ASM?	551
How long does a rescan take?	552
How can I control the bandwidth used for replication?	552
Are migrations performed by Application Migration Service crash consistent?	552
How can I perform an SSL connectivity and bandwidth test?	552
Why are MGN replication and conversion servers changing to Amazon Linux 2023?	554
AWS related	555
What does the AWS Transform MGN Machine Conversion Server do?	555
What boot modes are supported by the AWS Transform MGN?	556
How can we encrypt an unencrypted AWS Transform MGN base snapshot?	556
How do I change the server AMI on AWS after Migration?	557

Which AWS services are automatically installed when launching a test or cutover instance?	557
How long does it take to copy a disk from the AWS Transform MGN staging area to production?	557
What are the differences between conversion servers and replication servers?	557
Can I prevent AWS Transform MGN from cleaning up test instance resources in AWS?	558
Why are my Windows server disks read-only after launching the test or cutover instance?	558
What impacts the conversion and boot time of test and cutover instances?	558
Why do I observe EBS volume performance issues while using test or cutover instances? ..	559
What are the Amazon EBS volume limits for AWS Transform MGN?	559
How is the AWS Licensing Model Tenancy chosen for AWS Transform MGN?	560
How does AWS Transform MGN interact with Interface VPC Endpoints?	560
How do I use MGN with CloudWatch and EventBridge dashboards?	560
Does MGN work with...?	563
Does MGN work with Microsoft Windows Failover Clustering?	563
Does MGN work with Bitlocker encryption?	563
Does MGN work with FSx for ONTAP?	563
Post-launch actions related	564
What operating systems does the post-launch actions framework support?	564
What version of AWS Systems Manager Agent will be installed on my instance?	564
Why is the AWS Systems Manager Agent not executing my post launch actions?	564
Storage related	565
What target storage types does MGN support?	565
Can I use FSx for ONTAP for some servers and Amazon EBS for others?	565
Do I need to pre-install iSCSI packages on my source server?	566
What happens to FSx replication volumes after finalize cutover?	566
I initialized MGN before FSx for ONTAP support was available. Do I need to do anything?	566
Release notes	567
September 2026	567
August 2026	567
July 2026	568
June 2026	568
May 2026	569
April 2026	569

March 2026	569
December 2025	570
October 2025	570
September 2025	570
August 2025	570
July 2025	570
April 2025	571
February 2025	571
January 2025	571
October 2024	571
September 2024	571
August 2024	572
July 2024	572
June 2024	572
May 2024	572
March 2024	572
January 2024	573
December 2023	573
November 2023	573
September 2023	573
August 2023	574
July 2023	574
June 2023	574
May 2023	575
April 2023	575
March 2023	575
January 2023	576
November 2022	576
August 2022	577
July 2022	577
June 2022	577
May 2022	577
February 2022	577
January 2022	578
December 2021	578
November 2021	578

October 2021	578
July 2021	578
April 2021	578
Document history	579

What Is AWS Transform MGN?

AWS Transform MGN (MGN) automates the migration of physical, virtual, and cloud servers to AWS with minimal downtime, typically cutover windows of minutes. MGN performs continuous block-level replication of your source servers and converts them for launch on AWS, allowing you to migrate large numbers of servers without compatibility issues or performance disruption.

MGN works across a broad range of operating systems including Windows Server and various Linux distributions, and supports both IPv4 and IPv6 network configurations. You can replicate into standard Availability Zones or AWS Local Zones without any special configuration.

The service uses three configurable templates (replication, launch, and post-launch) to control how servers are replicated, launched, and configured after migration. Template settings are applied to each newly added server, and you can override settings for individual servers at any time.

To manage migrations at scale, you can group servers into applications and applications into waves. Configuration changes and actions such as launch, cutover, and archival can be performed at the server, application, or wave level, enabling bulk operations across large environments.

After migration, you can use AWS services to replatform or refactor your applications, making rehosting a fast first step toward modernization.

Accessing the AWS Transform MGN console

You can access AWS Transform MGN through the AWS Management Console or through the following link:

<https://console.aws.amazon.com/mgn/home>

Supported AWS Regions

The following AWS Regions are supported by AWS Transform MGN:

Region name	Region identity	Support in AWS Transform MGN
US East (Ohio)	us-east-2	Yes
US East (N. Virginia)	us-east-1	Yes

Region name	Region identity	Support in AWS Transform MGN
US West (N. California)	us-west-1	Yes
US West (Oregon)	us-west-2	Yes
Africa (Cape Town)	af-south-1	Yes
Asia Pacific (Hong Kong)	ap-east-1	Yes
Asia Pacific (Taipei)	ap-east-2	Yes
Asia Pacific (Thailand)	ap-southeast-7	Yes
Asia Pacific (Jakarta)	ap-southeast-3	Yes
Asia Pacific (Malaysia)	ap-southeast-5	Yes
Asia Pacific (New Zealand)	ap-southeast-6	Yes
Asia Pacific (Melbourne)	ap-southeast-4	Yes
Asia Pacific (Mumbai)	ap-south-1	Yes
Asia Pacific (Hyderabad)	ap-south-2	Yes
Asia Pacific (Osaka)	ap-northeast-3	Yes
Asia Pacific (Seoul)	ap-northeast-2	Yes
Asia Pacific (Singapore)	ap-southeast-1	Yes
Asia Pacific (Sydney)	ap-southeast-2	Yes
Asia Pacific (Tokyo)	ap-northeast-1	Yes
Canada (Central)	ca-central-1	Yes
Canada West (Calgary)	ca-west-1	Yes
Europe (Frankfurt)	eu-central-1	Yes

Region name	Region identity	Support in AWS Transform MGN
Europe (Zurich)	eu-central-2	Yes
Europe (Ireland)	eu-west-1	Yes
Europe (London)	eu-west-2	Yes
Europe (Paris)	eu-west-3	Yes
Europe (Milan)	eu-south-1	Yes
Europe (Spain)	eu-south-2	Yes
Europe (Stockholm)	eu-north-1	Yes
Israel (Tel Aviv)	il-central-1	Yes
Middle East (Bahrain)	me-south-1	Yes
Middle East (UAE)	me-central-1	Yes
Mexico (Central)	mx-central-1	Yes
South America (São Paulo)	sa-east-1	Yes
AWS GovCloud (US-East)	us-gov-east-1	Yes
AWS GovCloud (US-West)	us-gov-west-1	Yes

Learn more about [AWS Regional services](#).

AWS Transform MGN regional support includes [AWS Local Zones](#) associated with the above supported regions.

** Important**

MGN supports IPv6 in the regions where Amazon EC2 dual-stack endpoints are present. For the list of EC2 dual-stack endpoints, see [Dual-stack \(IPv4 and IPv6\) endpoints](#) in the *Amazon EC2 Developer Guide*.

MGN technical training materials

The following technical trainings are available for MGN:

- [AWS Transform MGN - A Technical Introduction](#)
- [Blog posts related to MGN](#)
- [MGN video playlist](#)

Quick start guide for AWS Transform MGN

Follow this guide to set up AWS Transform MGN for the first time, including:

Topics

- [First time setup](#)
- [Adding source servers](#)
- [Configuring launch settings](#)
- [Launching a test instance](#)
- [Launching a cutover instance](#)

First time setup

The first setup step for AWS Transform MGN is creating the replication template.

Choose **Get started** on the AWS Transform MGN landing page.

You will automatically be prompted to initialize the service the first time you log into AWS Transform MGN.

Initializing the service will create a replication template. This template will determine how data replication will work for each newly added source server.

The configured replication settings can be changed at any time for any individual source server or group of source servers. [Learn more about replication settings.](#)

[Learn more about changing individual server and multiple server replication settings.](#)

Important

Before configuring your replication template, ensure that you meet the [Network requirements for running AWS Transform MGN](#).

Once AWS Transform MGN is initialized you'll be redirected into the MGN console **Source servers** page.

To edit your replication template, choose **Replication template**. You will be able to edit individual server replication settings after adding your source servers to AWS Transform MGN.

The next step of the setup process is adding your source servers to AWS Transform MGN.

Adding source servers


Add source servers to AWS Transform MGN by installing the AWS Replication Agent (the Agent) on them. The Agent can be installed on both Linux and Windows servers. [Learn more about adding source servers.](#)

Note

If you are using the agentless replication for vCenter feature, then you will need to add your source servers by installing the MGN vCenter Client. [Learn more about agentless replication.](#)

Before adding your source servers, ensure that you meet all of the [network requirements](#).

The following is the MGN agent network architecture diagram:



Migration lifecycle

After the source server has been added to AWS Transform MGN, it will undergo the migration lifecycle steps.

The migration lifecycle shows the current state of each source server within the migration process. Lifecycle states include:

- **Not ready** – The server is undergoing the initial sync process and is not yet ready for testing. Data replication can only start once all of the initial sync steps have been completed.
- **Ready for testing** – The server has been successfully added to AWS Transform MGN and data replication has started. test or cutover instances can now be launched for this server.
- **Test in progress** – A Test instance is currently being launched for this server.
- **Ready for cutover** – This server has been tested and is now ready for a cutover instance to be launched.
- **Cutover in progress** – A cutover instance is currently being launched for this server.
- **Cutover complete** – This server has been cutover. All of the data on this server has been migrated to the AWS cutover instance.

- **Disconnected** – This server has been disconnected from AWS Transform MGN.

[Learn more about the migration lifecycle states.](#)

After the initial sync process has completed successfully, data replication will start automatically.

Configuring launch settings

After you have added your source servers to the AWS Transform MGN console, you will need to configure the launch settings for each server. The launch settings are a set of instructions that determine how a test or cutover instance will be launched for each source server on AWS. You must configure the launch settings before launching test or cutover instances. You can use the default settings or configure the settings to fit your requirements.

Note

You can change the launch settings after a test or cutover instance has been launched. You will need to launch a new test or cutover instance for the new settings to take effect.

You can access the launch settings by choosing the source server name of a source server on the **Source servers** page.

Within the individual server view, navigate to the **Launch settings** tab.

Here you can see your **General launch settings** and **EC2 launch template**. Choose the **Edit** button to edit your Launch settings or **Modify** to change your EC2 launch template.

Launch settings are composed of the following:

- **Instance type right-sizing** – The Instance type right-sizing feature allows AWS Transform MGN to launch a test or cutover instance type that best matches the hardware configuration of the source server. When activated, this feature overrides the instance type selected in the EC2 launch template.
- **Start instance upon launch** – Choose whether you want to start your test and cutover instances automatically upon launch or whether you want to start them manually through the Amazon EC2 Console.

- **Copy private IP** – Choose whether you want Application Migration Service to verify that the private IP used by the test or cutover instance matches the private IP used by the source server (not supported for IPv6).
- **Transfer server tags** – Choose whether you want AWS Transform MGN to transfer any user-configured custom tags from your source servers to your test or cutover instance.
- **OS Licensing** – Choose whether you want to Bring Your Own Licenses (BYOL) from the source server to the test or cutover instance.

AWS Transform MGN automatically creates an **EC2 launch template** for each new source server. AWS Transform MGN bases the majority of the instance launch settings on this template. You can edit this template to fit your needs.

[Learn more about Launch settings.](#)

Launching a test instance

After you have added all of your source servers and configured their launch settings, you are ready to launch a test instance. It is crucial to test the migration of your source servers to AWS before initiating a cutover to verify that your source servers function properly within the AWS environment.

Important

It is a best practice to perform a test at least two weeks before you plan to migrate your source servers. This time frame allows you to identify potential problems and solve them, before the actual cutover takes place. After launching test instances, use either SSH (Linux) or RDP (Windows) to connect to your instance and ensure that everything is working correctly.

You can test one source server at a time, or simultaneously test multiple source servers. For each source server, you will be informed of the success or failure of the test. You can test your source server as many times as you want. Each new test first deletes any previously launched test instance and dependent resources. Then, a new test instance is launched, which reflects the most up-to-date state of the source server. After the test, data replication continues as before. The new and modified data on the source server is transferred to the staging area subnet and not to the test instances that were launched during the test.

 Note

- When launching a test or cutover instance, you can launch up to 100 source servers in a single operation. Additional source servers can be launched in subsequent operations.
- Windows source servers need to have at least 2 GB of free space to successfully launch a test instance.
- Take into consideration that once a test instance is launched, actual resources will be used in your AWS account and you will be billed for these resources. You can terminate the operation of launched Test instances once you verify that they are working properly without impact to data replication.

Ready for testing indicators

Before launching a Test instance, ensure that your source servers are ready for testing by looking for the following indicators on the **Source servers** page:

1. Under the **Migration lifecycle** column, the server should show **Ready for testing**.
2. Under the **Data replication status** column, the server should show the **Healthy** status.
3. Under the **Next step** column, the server should show **Launch test instance**.

Starting a test

To launch a test instance for a single source server or multiple source servers, take the following steps:

- Go to the **Source servers** page and check the box to the left of each server for which you want to launch a test instance.
- Open the **Test and cutover** menu.
- Under **Testing**, choose the **Launch test instances** option to launch a test instance for this server.
- When the **Launch test instances for X servers** dialog appears, choose **Launch** to begin the test.

The AWS Transform MGN console will indicate **Launch job started** when the test has started.

Choose **View job details** on the dialog to view the specific Job for the test launch in the **Launch History** tab.

Successful test launch indicators

You can tell that the Test instance launch started successfully through several indicators on the **Source Servers** page.

1. The Alerts column will show the **Launched** status, indicating that a test instance has been launched for this server.
2. The **Migration lifecycle** column will show **Test in progress**.
3. The **Next step** column will show **Complete testing and mark as 'Ready for cutover'**.

Reverting or finalizing a test

After you have launched your test instances, open the Amazon EC2 Console and SSH or RDP into your test instances to ensure that they function correctly. Validate connectivity and perform acceptance tests for your application.

Reverting a test

If you encounter any issues and want to launch new test instances, or if you are performing a scheduled test and plan to perform additional tests before cutover, you can revert the test. This will revert your source servers' **Migration lifecycle** status to **Ready for testing**, indicating that these servers still require additional testing before they are ready for cutover. During a revert, you will also have the option to delete your Test instances for cost-saving purposes.

To revert a test:

1. Check the box to the left of every source server that has a launched test instance for which you want to revert the test.
2. Open the **Test and cutover** menu.
3. Under **Testing**, choose **Revert to "ready for testing"**
4. When the **Revert testing for X servers** dialog appears, select whether you want to terminate the launched instances used for testing. It is recommended to terminate these instances, as you will be charged for them even though you will no longer need them. Check the **Yes, terminate launched instances (recommended)** box and choose **Revert**.

The AWS Transform MGN console will indicate that testing has been reverted. The selected source servers' **Migration lifecycle** column will show the **Ready for testing** status, the **Next step** column will show **Launch test instance** and the launched Test instances will be deleted if that option was selected.

Marking as Ready for cutover

If you are completely done with your testing and are ready for cutover, you can finalize the test. This will change your source servers' **Migration lifecycle** status to **Ready for cutover**, indicating that all testing is complete and that these servers are now ready for cutover. You will also have the option to delete your Test instances for cost saving purposes.

To finalize a test:

1. Check the box to the left of every source server that has a launched Test instance for which you want to finalize the test.
2. Open the **Test and Cutover** menu.
3. Under **Testing**, choose **Mark as "Ready for cutover"**
4. Mark X servers as "Ready for cutover" dialog will appear. Select whether you want to terminate the launched instances used for testing. It is recommended to terminate these instances, as you will be charged for them even though you will no longer need them. Check the **Yes, terminate launched instances (recommended)** box and choose **Continue**.
5. The AWS Transform MGN console will confirm that the servers were marked as ready for cutover.

The console will indicate that testing has been finalized. The selected source servers' **Migration lifecycle** column will show the **Ready for cutover** status and the launched test instances will be deleted if that option was selected. The **Next step** column will show **Terminate launched instance; Launch cutover instance**.

6. You can now terminate the launched test instance directly from the Amazon EC2 Console as that instance is no longer needed (if you have not done so already through the MGN Console). You can quickly access the Test instance by navigating to the specific servers > **Server Details** > **Migration dashboard** > **Lifecycle** > **Launch status** and choosing **View in EC2 Console**.
7. The Amazon EC2 Console will automatically search for and display the test instance. Select the instance, open the **Instance state** menu, and choose **Terminate instance**. When the confirmation dialogue appears, choose **Terminate**.

Launching a cutover instance

Once you have finalized the testing of all of your source servers, you are ready for cutover. You should perform the cutover at a set date and time. The cutover will migrate your source servers to the cutover instances on AWS.

Important

It is a best practice to perform a test at least two weeks before you plan to migrate your source servers. This time frame allows you to identify potential problems and solve them, before the actual migration takes place. After launching test instances, use either SSH (Linux) or RDP (Windows) to connect to your instance and ensure that everything is working correctly.

You can cutover one source server at a time, or simultaneously cutover multiple source servers. For each source server, you will be informed of the success or failure of the cutover. For each new cutover, AWS Transform MGN first deletes any previously launched Test instance and dependent resources. Then, it launches a new cutover instance which reflects the most up-to-date state of the source server. After the cutover, data replication continues as before. The new and modified data on the source server is transferred to the Staging Area Subnet, and not to the cutover instances that were launched during the cutover.

Ready for cutover indicators

Before launching a cutover instance, ensure that your source servers are ready for cutover by looking for the following indicators on the **Source Servers** page:

1. Under the **Migration lifecycle** column, the server should show **Ready for cutover**.
2. Under the **Data replication status** column, the server should show the **Healthy** status.
3. Under the **Next step** column, the server should show **Terminate launched instance; Launch cutover instance** if you have not terminated your latest launched test instance.
4. Alternatively, the Next step column will show **Launch cutover instance** if you have terminated your latest launched test instance.

Starting a cutover

To launch a cutover instance for a single source server or multiple source servers, take the following steps:

1. Go to the **Source servers** page and check the box to the left of each server you want to cutover.
2. Open the **Test and cutover** menu.
3. Under **Cutover**, choose the **Launch cutover instances** option.
4. When the **Launch cutover instances for X servers** dialog appears, choose **Launch** to begin the cutover.

On the **Source servers** page, the **Migration lifecycle** column will show **Cutover in progress** and the **Next step** column will show **Finalize cutover**. When the cutover starts, the Application Migration Service Console will indicate **Launch job started**.

5. Choose **View job details** on the dialog to view the specific Job for the cutover launch in the **Launch History** tab.

Successful cutover launch indicators

You can tell that the cutover instance launch was started successfully through several indicators on the **Source servers** page.

1. The **Alerts** column will state **Launched**.
2. The **Migration lifecycle** column will state **Cutover in progress**.
3. The **Data replication status** will state **Healthy**.
4. The **Next step column** will state **Finalize cutover**.

Reverting or finalizing a cutover

Once you have launched your cutover instances, open the Amazon EC2 Console and SSH or RDP into your cutover instances to ensure that they function correctly. Validate connectivity and perform acceptance tests for your application.

Note

You should turn on Termination Protection after you have completed your testing and before you are ready to finalize the cutover. Learn more about enabling termination protection in [this Amazon EC2 article](#).

Reverting a cutover

If you encounter any issues and want to launch new cutover instances, you can revert the cutover. This will revert your source servers' **Migration lifecycle** status to **Ready for cutover**, indicating that these servers have not undergone cutover. During a revert, you will also have the option to delete your Cutover instances for cost-saving purposes.

To revert a cutover:

1. Check the box to the left of every source server that has a launched cutover instance you want to revert.
2. Open the **Test and cutover** menu.
3. Under **Cutover**, choose **Revert to "ready for cutover"**
4. This will revert your source servers' **Migration lifecycle** status to **Ready for cutover**, indicating that these servers have not undergone cutover.

When the **Revert cutover for X servers** dialog appears, choose **Revert**.

Finalizing a cutover

If you are completely done with your migration and performed a successful cutover, you can finalize the cutover. This will change your source servers' **Migration lifecycle** status to **Cutover complete**, indicating that the cutover is complete and that the migration has been performed successfully. In addition, this will stop data replication and cause all replicated data to be discarded. All AWS resources used for data replication will be terminated.

To finalize a cutover:

1. Check the box to the left of every source server that has a launched cutover instance you want to finalize.
2. Open the **Test and cutover** menu.

3. Under **Cutover**, choose **Finalize cutover**.
4. The **Finalize cutover for X servers** dialog will appear. Choose **Finalize**. This will change your source servers' **Migration lifecycle** status to **Cutover complete**, indicating that the cutover is complete and that the migration has been performed successfully. In addition, this will stop data replication and cause all replicated data to be discarded. All AWS resources used for data replication will be terminated.

The AWS Transform MGN console will indicate **Cutover finalized** when the cutover has completed successfully.

The AWS Transform MGN console will automatically stop data replication for the source servers that were cutover to save resource costs. The selected source servers' **Migration lifecycle** column will show the **Cutover complete** status, the **Data replication** status column will show **Disconnected**, and the **Next step** column will show **Mark as archived**. The source servers have now been successfully migrated into AWS.

5. You can now archive your source servers that have launched cutover instances. Archiving will remove these source servers from the main **Source servers** page, allowing you to focus on source servers that have not yet been cutover. You will still be able to access the archived servers through filtering options.

To archive your cutover source servers:

- a. Check the box to the left of each source server for which the **Migration lifecycle** column states **Cutover complete**.
- b. Open the **Actions** menu and choose **Mark as archived**.
- c. When the **Archive X server** dialog appears, choose **Archive**.
- d. To see your archived servers, choose **Archived source servers** from the drop-down menu in the source servers view.

You will now be able to see all of your archived servers. Use the same drop-down menu to see only **Active source servers** or **Discovered source servers**, according to your preferences.

Getting started with AWS Transform MGN

Learn how to get started with AWS Transform MGN. Review the migration workflow, learn how to initialize MGN in either the console or using the API, and review best practices.

Topics

- [Migration workflow](#)
- [Initializing MGN with the console](#)
- [Initializing AWS Transform MGN with the API](#)
- [Configuring the templates](#)
- [Using the AWS Transform MGN console](#)
- [Best practices for AWS Transform MGN](#)
- [Migration at scale using AWS Transform MGN](#)
- [AWS Transform MGN service quota limits](#)

Migration workflow

The general process is:

1. Initialize AWS Transform MGN in the target region. Refer to the [list](#) of supported AWS regions.
2. Install the AWS Replication Agent on the source server. Learn more about [agent installation](#).

Note

If you are using the agentless replication for vCenter feature, then you will need to add your source servers by installing the MGN vCenter Client. [Learn more about agentless replication.](#)

3. Wait until the initial sync is finished. After installing the agent, you need to wait for the initial synchronization process to complete. This process performs block level replication from the source server to the replication server in staging area.
4. Launch test instances. Once the initial sync is finished, you can launch a target machine in Test Mode. This allows you to perform acceptance testing and verify that the migrated environment is functioning correctly.

5. Perform acceptance tests on the servers. After the test instance is tested successfully, finalize the test and delete the test instance.
6. Configure Post-launch actions (if needed). Learn more about [Post-launch settings](#).
7. Wait for the cutover window.
8. Confirm that there is no lag.
9. Stop all operational services on the source server.
10. Launch a cutover instance. Launch the target machine in Cutover Mode, which initiates the final migration process.
11. Confirm that the cutover instance was launched successfully and then finalize the cutover.
12. Archive the source server.

Initializing MGN with the console

In order to use AWS Transform MGN (MGN), the service must first be initialized for any AWS Region in which you plan to use MGN.

You can initialize the service via the console or via the API.

During the initialization process:

- The required IAM roles and policies will be created.
- The required templates are configured.

AWS Transform MGN must be initialized upon first use from within the MGN console by creating a replication template.

Once you create the replication template, the initialization process takes place automatically.

Note

The AWS Transform MGN can only be initialized by the IAM user with the "AdministratorAccess" managed policy attached in your AWS account.

For information on the IAM roles that MGN creates on your behalf during the initialization process, see [IAM role creation](#). For information on the predefined managed IAM policies that MGN includes, see [Additional policies](#).

You can also initialize MGN using the API. For more information, see [Initializing AWS Transform MGN with the API](#).

IAM role creation

During initialization the following IAM roles will be created.

1. **AWSServiceRoleForApplicationMigrationService**
2. **AWSApplicationMigrationReplicationServerRole**
3. **AWSApplicationMigrationConversionServerRole**
4. **AWSApplicationMigrationMGHRole**
5. **AWSApplicationMigrationLaunchInstanceWithDrsRole**
6. **AWSApplicationMigrationLaunchInstanceWithSsmRole**
7. **AWSApplicationMigrationAgentRole**

Learn more about [AWS Transform MGN roles and managed policies](#).

Additional policies

You can create roles with granular permission for AWS Transform MGN. The service comes with the following predefined managed IAM policies:

- **AWSApplicationMigrationFullAccess** – This policy provides permissions to all public APIs of AWS Transform MGN (AWS MGN), as well as permissions to read AWS KMS key information.
- **AWSApplicationMigrationEC2Access** – This policy allows Amazon EC2 operations required to use AWS Transform MGN (AWS MGN) to launch the migrated servers as Amazon EC2 instances.
- **AWSApplicationMigrationSSMAccess** – This policy allows Amazon EC2 Systems Manager operations required to use AWS Transform MGN (AWS MGN) to run SSM documents post migration of source servers.
- **AWSApplicationMigrationReadOnlyAccess** – The read-only policy allows the user to view all data available in the AWS MGN console but does not allow them to modify any data or perform any actions. This policy also includes several Amazon EC2 read-only permissions.
- **AWSApplicationMigrationAgentPolicy** – This policy allows a user to install the AWS Replication Agent. [Learn more about installing the AWS Replication Agent](#).
- **AWSApplicationMigrationAgentInstallationPolicy** – This policy allows a user to install the AWS Replication Agent. [Learn more about installing the AWS Replication Agent](#).

- **AWSApplicationMigrationServiceEc2InstancePolicy** – This policy allows installing and using the AWS Replication Agent, which is used by MGN (AWS MGN) to migrate source servers that run on Amazon EC2 (cross-Region or cross-AZ). An IAM role with this policy should be attached (as an Amazon EC2 Instance Profile) to the Amazon EC2 Instances.

You can find all of these policies in the [IAM Console](#).

 Note

You must attach the `AWSApplicationMigrationFullAccess` and the `AWSApplicationMigrationEC2Access` policies to your users and roles to be able to launch test and cutover instances and to complete a full migration cycle with AWS MGN.

Initializing AWS Transform MGN with the API

In order to use AWS Transform MGN (MGN), the service must first be initialized for any AWS Region in which you plan to use MGN.

You can initialize the service via the console or via the API.

During the initialization process:

- The required IAM roles and policies will be created.
- The required templates are configured.

You can initialize AWS Transform MGN through the API. This option allows you to automate service initialization through a script when initializing multiple accounts.

You can also initialize MGN using the console. For more information, see [Initializing MGN with the console](#).

To initialize the service via the API, take the following steps:

1. Create the required IAM roles.
2. Create the replication template and launch template.

Note

You must complete both steps to finalize the service initialization process.

Creating the required IAM roles

To initialize MGN with the API, create the following IAM roles through the [IAM CreateRole API](#). Learn more about [creating IAM roles in the AWS IAM documentation](#). Creation of each role must include the following parameters:

Role name	Trusted entities		
	Principal	Action	Condition
AWSApplicationMigrationReplicationServerRole	"ec2.amazonaws.com"	"sts:AssumeRole"	-
AWSApplicationMigrationConversionServerRole	"ec2.amazonaws.com"	"sts:AssumeRole"	-
AWSApplicationMigrationMGHRole	"mgn.amazonaws.com"	"sts:AssumeRole"	-
AWSApplicationMigrationLaunchInstanceWithDrsRole	"ec2.amazonaws.com"	"sts:AssumeRole"	-
AWSApplicationMigrationLaunchInstanceWithSsmRole	"ec2.amazonaws.com"	"sts:AssumeRole"	-
AWSApplicationMigrationFsxProxyRole	"mgn.amazonaws.com"	"sts:AssumeRole"	-
AWSApplicationMigrationFsxProxyLinkRole	"mgn.amazonaws.com"	"sts:AssumeRole"	-

Role name	Trusted entities		
	Principal	Action	Condition
AWSApplicationMigrationAgentRole	"mgn.amazonaws.com"	["sts:AssumeRole", "sts:SetSourceIdentity"]	{ "StringLike": { "sts:SourceIdentity": "s-*", "aws:SourceAccount": "<SOURCE-ACCOUNT-ID>" } }

1. Attach Managed Policy **AWSApplicationMigrationReplicationServerPolicy** to Role **AWSApplicationMigrationReplicationServerRole**
2. Attach Managed Policy **AWSApplicationMigrationConversionServerPolicy** to Role **AWSApplicationMigrationConversionServerRole**
3. Attach Managed Policy **AWSApplicationMigrationMGHAccess** to Role **AWSApplicationMigrationMGHRole**
4. Attach Managed Policies **AmazonSSMManagedInstanceCore** and **AWSElasticDisasterRecoveryEc2InstancePolicy** to Role **AWSApplicationMigrationLaunchInstanceWithDrsRole**
5. Attach Managed Policy **AmazonSSMManagedInstanceCore** to Role **AWSApplicationMigrationLaunchInstanceWithSsmRole**
6. Attach Managed Policy **AWSApplicationMigrationFSxProxyPolicy** to Role **AWSApplicationMigrationFsxProxyRole**
7. Attach Managed Policy **AWSApplicationMigrationFSxProxyVPCPolicy** to Role **AWSApplicationMigrationFsxProxyLinkRole**
8. Attach Managed Policy **AWSApplicationMigrationAgentPolicy_v2** to Role **AWSApplicationMigrationAgentRole**

Once the policies are attached to the roles, run the `aws mgn initialize-service` command. This will automatically create the service-linked role, create instance profiles, and add Roles to Instance Profiles. After running this command, you must still create the replication configuration template and launch configuration template to finalize initialization.

[Learn more about AWS Transform MGN roles and managed policies.](#)

Creating the templates

To finalize the initialization process, you will need to [create the replication template](#) and launch template by running the following commands:

- `aws mgn create-replication-configuration-template`
- `aws mgn create-launch-configuration-template`

Configuring the templates

As part of the initialization of AWS Transform MGN, you have the opportunity to configure three templates.

- Replication template (mandatory)
- Launch template (optional)
- Post-launch template (optional)

Configuring your replication template

To initialize AWS Transform MGN, you must first create and configure a replication template. If you initialize MGN with the console, the initialization process creates the template for you. If you initialize MGN with the API, you create the template. For more information, see [Initializing AWS Transform MGN with the API](#)

The replication template determines how data replication will work for each new server you add. The settings configured in this template will be applied to each newly added source server. [Learn more about the replication template.](#)

Configuring your launch template

As part of the AWS MGN initialization process, you can configure your launch template. Every source server added to AWS MGN has launch settings that control actions performed after the server is launched in AWS. These settings are created automatically based on this default launch template, which can be modified at any time. You can also choose to modify the launch template for an individual source server.

[Learn more about the launch template.](#)

Configuring a post-launch template

As part of the AWS MGN initialization process, you can configure your post-launch template. The post-launch template controls which post-launch actions will be executed when launching new instances. These settings are created automatically for each server based on the post-launch template and can be modified at any time. You can also modify the post-launch settings for any individual source server.

To configure the post-launch actions, complete the following steps:

1. [Activate post-launch actions.](#)
2. [Configure predefined post-launch actions](#) according to your preferences.
3. [Create custom post-launch actions](#) according to your preferences.

Using the AWS Transform MGN console

AWS Transform MGN is AWS Region-specific. Ensure that you select the correct Region from the **Select a Region** menu when using AWS Transform MGN, just like you would with other AWS Region-specific services such as Amazon EC2.

AWS Transform MGN is divided into several primary pages. Each page contains additional tabs and actions. The default view for the AWS Transform MGN console is the **Source servers** page. This page automatically opens every time you open AWS Transform MGN.

You can navigate to other AWS Transform MGN pages through the left-hand **AWS Transform MGN** navigation menu. You can also access the AWS Transform MGN documentation through this menu.

Each AWS Transform MGN page will open in the right-hand main view. Here, you can interact with the various tabs, actions, and settings on the page.

Source servers page

The **Source servers** page lists all of the source servers you added to AWS Transform MGN and allows you to interact with your servers and perform a variety of actions. [Learn more about the source servers page.](#)

- To control your source servers within the AWS Transform MGN console, use the **Actions**, **Replication**, and **Test and cutover** menus. The commands within those menus influence the specific source servers you have selected. You can select a single source server or multiple source servers for any command by checking the box to the left of the server name.
- To review the progress of all commands, use the **Launch history** tab. [Learn more about Launch history.](#)
- Use the **Filter source servers...** box to filter servers based on a variety of parameters.

AWS Transform MGN color codes the state of each source server. Use the **Alerts** column to easily determine the state of your server.

A clock icon with gray text indicates that the server is experiencing temporary issues. The nature of the issue will be identified (for example, "Lagging").

A red x and text indicates that the server is experiencing significant issues that need to be addressed before replication can resume. The nature of the issue will be indicated (for example, "Stalled")

When various commands are initiated, AWS Transform MGN will display information messages at the top of the **Source Servers** page. AWS Transform MGN color codes these messages for clarity.

- A green message indicates that a command was completed successfully.
- A red message means that a command was not completed successfully.

Each message shows details and links to supplemental information.

AWS Transform MGN allows you to interact with and manage each server. Choose the source server name to be redirected to the Server details view.

The **Server details** tab shows specific details for an individual server. From here, you can review the server's migration lifecycle and health, view various technical details, manage tags, manage disks, edit the server's replication settings, and edit the server's launch settings through the various tabs. [Learn more about the Server details view.](#)

Certain AWS Transform MGN commands, such as **Edit replication settings**, allow you to interact with multiple source servers at once. When multiple source servers are selected by checking the box to the left of the server name and the **Replication > Edit replication settings** option is chosen, AWS Transform MGN will indicate which servers are being edited.

For your changes to take effect, choose **Save** at the bottom of each settings page.

Best practices for AWS Transform MGN

Planning

During the phases of your migration project, review these best practices to help you to a successful outcome.

1. Plan your Migration project before installing the AWS Replication Agent on your source servers.
2. Do not perform any reboots on the source servers before a cutover.
3. Do not archive or disconnect the source server from AWS until your launched cutover instance in AWS is working as expected.

Testing

1. Perform a test at least two weeks before you plan to migrate your source servers. This time frame is intended for identifying potential problems and solving them, before the actual cutover takes place. After performing the test launch, validate connectivity to your test instances (using SSH for Linux or RDP for Windows), and perform acceptance tests for your application.
2. Ensure that you perform a Test before performing a cutover.

Successful implementation

The following are the required steps to complete a successful migration implementation with AWS Transform MGN:

1. Deploy the AWS Replication Agent on your source servers.
2. Confirm that the data replication status is **Healthy**.
3. Test the launch of Test instances a week before the actual cutover.
4. Address any issues that come up, such as Launch setting misconfiguration and potential AWS limits.
5. Launch cutover instances for the servers on the planned date.

Ensuring project success

1. Train a field technical team & assign an AWS Transform MGN SME.
2. Share project timelines with AWS Transform MGN.
3. Monitor data replication progress and report any issues in advance.
4. Perform a test for every server in advance, and report issues to AWS Transform MGN.
5. Coordinate cutover windows with AWS Transform MGN in advance.

Migration at scale using AWS Transform MGN

Choose your agent-based vs agentless migration approach

When planning your migration with AWS Transform MGN (MGN), review which approach best suits your environment and goals. MGN offers agentless and agent-based replication to migrate your workloads into AWS EC2. Agent-based replication is supported for both VMware and non-VMware environments, while agentless replication is supported only for VMware environments.

1. Consider agentless replication if your security policies restrict installing agents at the OS level of your Virtual Machines.
2. Review the network requirements for both, [agentless replication](#) and [agent-based replication](#), to determine the best approach for your needs.
3. Before beginning your migration, ensure your planned operating systems are [supported](#) by your chosen replication method.
4. We recommend reviewing [AWS Prescriptive Guidance](#) for a rehost migration approach with MGN to ensure it meets your business needs.

Cost planning

Before migrating, review the associated costs of using the service. For full pricing details, see the [MGN pricing page](#).

1. For each migrated server, MGN can be used free of MGN service charge, for 2,160 hours, which is 90 days when used continuously. Resources that are provisioned in the replication and migration processes, such as EC2, Amazon EBS volumes, and FSx for ONTAP storage, incur charges according to your account rates.

2. Servers with high-volume workloads may require higher-performance replication resources, potentially increasing costs. Conduct network and storage benchmarking to determine the necessary staging area resources and estimate associated expenses.
3. Plan your software licensing strategy when migrating servers. Discuss licensing terms with your operating system (OS) and software vendors as needed.
4. When servers fail to launch, the Support team may request permission to run diagnostics. This iterative process helps the service team address edge cases causing launch failures. Diagnostic launches will incur charges based on the source and conversion server settings. Begin testing at least two weeks before your migration deadline. This allows time to resolve unexpected issues before the final cutover.

Import/Export

When performing large-scale migrations, you can use the [Import/Export](#) feature in AWS Transform MGN to efficiently stage your servers with the correct configurations. This feature uses a CSV template file with predefined columns that map to attributes of Waves, Applications, and Source Servers in MGN. The template includes columns for Launch Template settings, enabling bulk edits of these configuration options. Additionally, you can specify Tags to assist with cost center association, ensuring your cost planning remains accurate. By using the Import/Export feature, you can streamline your migration process and maintain better control over your server configurations and associated costs.

MGN Connector

When using AWS Transform MGN (MGN), first verify your servers meet the prerequisites for agent installation, then install the agent on all source servers. The [MGN Connector](#) assists with these tasks. Before using the MGN Connector, verify its prerequisites and ensure the MGN Connector can communicate with source servers via SSH for Linux and WinRM for Windows to execute commands.

Network benchmarking

Network benchmarking is crucial for estimating data replication performance from a source environment to Amazon Web Services (AWS). The time required to replicate data depends on various factors, with network bandwidth performance being a key consideration. To accurately assess expected performance, we recommend testing the network bandwidth between the source environment and AWS.

1. For this assessment, we recommend using the SSL connectivity and bandwidth test Amazon Machine Image (AMI). This AMI is preferable to alternatives like iperf because it performs the test with encryption, accurately simulating the replication agent's behavior. Detailed instructions for conducting this test are available for [Windows](#) and [Linux](#) servers.
2. To determine whether the bandwidth test results are sufficient for your intended workloads, compare them with the storage benchmarking results discussed in the subsequent section. This comparison will provide a comprehensive view of your system's replication capabilities. After completing the network bandwidth test, proceed to the storage benchmarking section to finalize your performance assessment. Based on these combined results, you can determine if any network or storage optimizations are necessary before initiating data replication. This thorough evaluation ensures that your infrastructure is adequately prepared for efficient data transfer to AWS.

Storage benchmarking

Performing storage benchmarking, in addition to network benchmarking, helps understand MGN replication resource provisioning and minimize performance issues. By understanding the total amount of storage and the daily change rate per server, you can determine if the network bandwidth will be sufficient for data replication with MGN. To capture storage performance metrics, use [iostat](#) on Linux, and [Performance Monitor](#) on Windows. These tools will help you determine if the network bandwidth can keep up with the Writes/s on your disks. Additionally, these tests will reveal if the source storage device is encountering any bottlenecks that could affect replication performance.

Dedicated replication servers

After performing your storage benchmarking, you will have a clear understanding of which servers are busy and have a high rate of data change. In these cases, it may be best to set these servers to use [Dedicated Replication server](#).

1. This approach allows you to select an EC2 Instance type that can handle the IOPs and throughput from your source server. For a list of instance types and their EBS performance capabilities, refer to [Amazon EBS-optimized instance types](#). When selecting an instance type, consider your cost planning, as this will incur charges based on the instance type selected. For pricing details, refer to [Amazon EC2 On-Demand Pricing](#).
2. To achieve optimum replication resource performance and cost savings, it is recommended to use [AWS Compute Optimizer](#). This service helps identify over or under provisioned EC2 and EBS

resources, allowing you to optimize your setup. Before getting started with Compute Optimizer, review its [pricing](#) to ensure it aligns with your budget and needs.

Wave planning

Planning your migration in waves is crucial for successful large-scale migrations. With AWS Transform MGN you can group source servers and applications into waves, allowing you to migrate servers in batches rather than all at once.

1. Migrating servers in [waves](#) offers several benefits. It allows you to focus on a specific set of servers at a time and ensure the first wave operates as expected before proceeding to the next.
2. Grouping interdependent servers into [applications](#) helps prevent issues during migration. For example, migrating member servers before domain controllers could cause authentication problems and other Active Directory-related issues. This approach enables a more controlled and reliable migration process, reducing the risk of widespread disruptions and allowing for adjustments between waves as needed.

Cutover planning

Planning for your cutover date is crucial for a successful migration. Work backwards from this date to schedule tests and address potential issues, ensuring you meet your migration timelines.

1. Conduct comprehensive [testing](#) at least two weeks before your planned migration date. This timeframe allows you to resolve or escalate specific problems with [AWS Support](#). For legacy operating systems or unique server configurations, allocate more than two weeks for testing, as complex issues may require additional time to resolve.
2. If you need expert support during migration, consider using [AWS Countdown](#). This service provides designated engineers from a team of AWS specialists who offer proactive guidance and troubleshooting assistance throughout your migration process.
3. Train your team for AWS Cloud before migration to prepare for the change in the environment as upskilling your team is an important step in the migration journey.

Clustered servers

Migrating clustered servers to AWS requires careful planning due to its complexity. The AWS Transform MGN operates at the server level and is not cluster-aware, necessitating additional

steps after cutover to ensure proper functionality. Consider the following high-level approach for migrating clusters with MGN:

1. Begin by migrating the primary node to AWS. This establishes the foundation for your cluster in the new environment.
2. For secondary nodes, you have several options. You can use MGN to migrate these nodes after the primary node is operational in AWS. Alternatively, consider performing a native backup and restore, using log shipping, or employing other methods that suit your specific requirements.
3. Once all nodes are in AWS, recreate the cluster configuration. This step is crucial for restoring the cluster's functionality in the new environment. Pay special attention to clusters with shared storage. In these cases, shared drives will appear as local disks on the cutover EC2 Instance. Depending on your migration plan, it may be more effective to exclude shared drives from MGN replication. Instead, consider migrating this data to an AWS shared-storage solution such as Amazon FSx.
4. MGN's lack of cluster awareness means post-cutover steps are often necessary to ensure your clustered servers work as expected. Tailor your migration approach based on your specific cluster configuration and requirements to achieve the best results.

Environment updates/changes

We recommend maintaining a stable environment between the test launch and cutover launch to ensure a smooth migration. However, we recognize that critical updates may be necessary. In such cases, re-test before your cutover date to prevent unexpected issues during migration. The following changes could affect a migration with MGN:

1. Operating System updates or patches.
2. Server configuration modifications.
3. Network alterations.
4. Security software updates.

API limits

API limits can cause migration delays during large-scale migrations. It's crucial to review the limits of all services used in AWS migrations, including [MGN](#), [EC2](#), [EBS](#), [FSx for ONTAP](#), and [VPC](#).

1. The default limit for MGN concurrent replicating servers ("Max Active Source Servers") is 150. Increasing this limit requires a detailed justification and migration plan. Contact [AWS Support](#) if you need to request a limit increase.
2. EC2 and EBS service quotas will be based on your AWS Environment. Some customers may already have workloads running in AWS, while others may have none. Plan for the number of servers and disks being migrated to estimate how many EC2 instances will run in your account at a given time. This includes MGN Replication Server EC2 instances and their associated storage volumes (Amazon EBS or FSx for ONTAP). Additionally, account for the [EBS Snapshots](#) that MGN will create during the replication process.
3. If you plan to use a single IP address for network traffic from the [staging area](#) to the MGN endpoints (such as a security appliance or firewall), consult your AWS account team or AWS Support. Migrating in waves can help mitigate API limits or throttling issues.

It can be challenging to anticipate all AWS service limitations when planning your migration. Testing with the same number of servers you intend to migrate will help identify potential limitations. Developing a comprehensive test and cutover plan will enable you to address these issues promptly, ensuring a smooth migration experience. By following these guidelines and planning carefully, you can effectively navigate API limits and service quotas during your AWS migration.

AWS Transform MGN service quota limits

The following are the MGN service quota limits:

Name	Default	Description
Concurrent jobs in progress	Each supported AWS Region: 20	Launching a test or cutover instance, or a cleanup action is considered a "job". Multiple servers launched together count as a single job. This parameter is the maximum number of Jobs that can be run concurrently. Jobs that are "completed" are not counted against this quota.

Name	Default	Description
Max active source servers	Each supported AWS Region: 150	The maximum number of servers that can be actively replicating at any time. For larger migrations, contact Support. Learn more about requesting a quota increase.
Max non-archived source servers	Each supported AWS Region: 4,000	This parameter is used for agentless migrations. This is the max number of servers that can be managed by MGN, in non-archived state. This includes the servers that are actively replicating, as well as any servers whose replication has not yet started. The number of actively replicating servers is controlled by the parameter "Max active source servers". For larger migrations, contact Support. Learn more about requesting a quota increase
Max source servers in a single job	Each supported AWS Region: 200	Launching a test or cutover instance, or a cleanup action is considered a "Job". If you select multiple servers, and perform one of these actions, they are grouped into a single Job. This is the maximum number of servers that can be grouped into a single Job.

Name	Default	Description
Max source servers in all jobs	Each supported AWS Region: 200	Launching a test or cutover instance, or a cleanup action is considered a "job". This is the maximum total number of servers that can be configured in all active jobs. Jobs that are "completed" are not counted against this quota.
Max total source servers per AWS account	Each supported AWS Region: 50,000	This parameter is the maximum total servers, both active and archived, that can be migrated in a single account in each AWS Region. Servers that are deleted, are not counted against this quota.
Max concurrent jobs per source server	Each supported AWS Region: 1	Launching a test or cutover instance, or a cleanup action is considered a "job". This is the maximum number of active jobs, that can be configured per server. Jobs that are "completed" are not counted against this quota.
Max actions per source server	Each supported AWS Region: 100	This parameter is the maximum amount of custom post-migration actions that can be associated with a specific source server.

Name	Default	Description
Max actions per account template	Each supported AWS Region: 100	This parameter is the maximum amount of custom post-migration actions that can be associated with an account template.
Max source servers per application	Each supported AWS Region: 200	This parameter is the maximum total servers, both active and archived, that can be added to an application.
Max active applications	Each supported AWS Region: 200	This parameter is the maximum number of applications that are in an active status.
Max archived applications	Each supported AWS Region: 10,000	This parameter is the maximum number of applications that are in an archived status.
Max applications per wave	Each supported AWS Region: 1000	This parameter is the maximum total applications, both active and archived, that can be added to a wave.
Max active waves	Each supported AWS Region: 200	This parameter is the maximum number of waves that are in an active status.
Max archived waves	Each supported AWS Region: 10,000	This parameter is the maximum number of waves that are in an archived status.

Name	Default	Description
Global View - Max Managed Accounts	5,000	This is the maximum number of member accounts that can be managed by a single management account.

You can learn about the MGN service limits in the [AWS General Reference](#)

Configuring AWS Transform MGN Settings

AWS Transform MGN uses replication settings to determine how data is replicated from source servers to your AWS account and Region. Learn how to configure your initial replication template and how to set individual server replication settings.

You must configure the replication template upon first use of AWS Transform MGN. The replication template determines how your servers are replicated to AWS through settings such as Replication Server instance type, target storage type, security groups, data routing, and tags. The settings configured in the replication template are automatically used for every server you add to AWS Transform MGN.

Once you have configured your Replication template, you can make changes to individual servers or a group of servers by editing their replication settings within the Server Details View.

You can also configure optional post-launch settings that automate target instance deployment and prepare your migrated servers for disaster recovery with AWS Elastic Disaster Recovery.

Topics

- [Replication template](#)
- [Launch template](#)
- [Post-launch template](#)

Replication template

The replication template determines how data is replicated from your source servers to AWS. Configure the replication template before adding source servers to AWS Transform MGN. You can modify the template at any time. Template settings are applied to each newly added server, but changes do not affect servers already in the MGN console.

You can also edit the replication settings for a particular server or group of servers after you add them to AWS Transform MGN. You can also control other source server settings through the **Settings** section in the menu.

Topics

- [Understanding template settings and server-specific settings](#)

- [Edit your replication settings template](#)
- [Edit replication settings for a server](#)
- [Replication template reference](#)

Understanding template settings and server-specific settings

The replication template determines how data replication works for each new server you add to AWS Transform MGN. You are prompted to configure your replication template upon your first use of AWS Transform MGN. These settings are applied to each source server you add, but are not applied to servers already in the console when the template changes are made.

For servers already in the console, you can change the replication settings for individual source servers or for a group of source servers. These changes do not affect the replication settings template. [Learn more about configuring your initial replication template.](#)

Edit your replication settings template

To edit the replication settings template:

- Choose **Replication template**, under **Settings** in the console menu.
- This opens the **Replication template** view. Choose **Edit** to edit your account-wide replication settings. These changes are applied to each server you add to your account but do not affect servers that you already added to AWS Transform MGN.

Edit replication settings for a server

To edit the settings for an individual server or group of servers:

- Select servers on the **Source servers** page.
- Open the **Replication** menu and choose **Edit replication settings**.

The names of the servers you are editing appear under the **Selected servers** dropdown.

- Edit individual replication settings under the **Replication settings** category.
- To change the settings, choose the preferred option from the drop-down menu under each setting category.

Any setting that you don't change is labeled **Do not change** option

- Choose **Save replication settings**.

The replication settings categories are explained in the following sections.

Replication template reference

Replication servers are lightweight Amazon EC2 instances that are used to replicate data between your source servers and AWS. Replication servers are automatically launched and terminated as needed. You can modify the behavior of the replication servers by modifying the settings for a single source server or multiple source servers. Alternatively, you can run MGN with the default replication settings.

The replication options, include:

- The subnet within which the replication server is launched. The subnet can use both IPv4 and IPv6 CIDRs.
- Replication Server instance type
- Target storage type (Amazon EBS or FSx for ONTAP)
- Security groups
- IP Version - you can choose IPv4 or IPv6.

Staging area subnet

Choose the **Staging area subnet** that you want to allocate as the staging area subnet for all of your replication servers. If you chose to use IPv6, you must choose a subnet that uses IPv6 CIDRs.

The best practice is to create a single dedicated, separate subnet for all of your migration waves using your AWS account. Learn more about creating subnets in [this AWS VPC article](#).

If a default subnet does not exist, select a specific subnet. The drop-down menu contains a list of all subnets that are available in the console's AWS Region.

Note

Changing the subnet does not significantly interfere with ongoing data replication, although there may be a minor delay of several minutes while the servers are moved from one subnet to another.

Using multiple subnets

The best practice is to use a single staging area subnet for all of your migration waves within a single AWS account. You may want to use multiple subnets in certain cases, such as the migration of thousands of servers.

Note

Using more than one staging area subnet might result in higher compute consumption as more replication servers are needed.

Launching replication servers in Availability Zones

If you want your replication servers to be launched in a specific Availability Zone, then select or create a subnet in that specific Availability Zone. Learn more about using Availability Zones in [this Amazon Elastic Compute Cloud article](#).

Replication server instance type

Choose the **Replication server instance type**. This determines the instance type and size that is used for the launch of each replication server.

The best practice is to not change the default replication server instance type unless there is a business need for doing so. By default, AWS Transform MGN uses the **t3.small** instance type. This is the most cost effective instance type and should work well for most common workloads. You can change the replication server instance type to speed up the initial sync of data from your source servers to AWS. Changing the instance type will likely lead to increased compute costs.

You can choose the **Replication server instance type** from the drop-down menu, which contains all available types. Recommended and commonly used instance types are displayed first. You can also search for a specific instance type in the search box.

You can change the replication server instance type for servers that are replicating too slowly or servers that are constantly busy or experience frequent spikes. These are the most common instance type changes:

- Servers with less than 26 disks – Change the instance type to **m5.large**. Increase the instance type to m5.xlarge or higher as needed.

- Servers with more than 26 disks, or servers in AWS Regions that do not support m5 instance types – change the instance type to **m4.large**. Increase to m4.xlarge or higher, as needed.

Note

- Changing the replication server instance type will not affect data replication. Data replication will automatically continue from where it left off, using the new instance type you selected.
- By default, replication servers are automatically assigned a public IP address from Amazon's public IP space.
- Replication Servers are only supported on x86_64 CPU architecture instance types.

Dedicated instance for replication server

Choose whether you would like to use a **Dedicated instance for replication server**.

When an external server is very write-intensive, the replication of data from its disks to a shared Replication Server can interfere with the data replication of other servers. In these cases you should choose the **Use dedicated replication server** option (and also consider changing Replication server instance type).

Otherwise, choose the **Do not use dedicated replication server** option.

Note

Using a dedicated replication server may increase the Amazon EC2 cost you incur during replication.

Target storage type

AWS Transform MGN supports two target storage types for replication. You can select the storage type that best suits your workload requirements:

- **Amazon Elastic Block Store (Amazon EBS)** – The default storage type. Amazon EBS volumes are used by the replication servers. For configuration details, see [Amazon EBS configuration](#).

- **Amazon FSx for NetApp ONTAP (FSx for ONTAP)** – Enterprise file storage with NetApp ONTAP capabilities. When selecting FSx for ONTAP, you must provide an FSx Storage Secret ARN and a security group configured for iSCSI access. For setup instructions, see [FSx for ONTAP configuration](#).

Store snapshots in AWS Local Zone

MGN does not require special configuration for Local Zones. If you select a subnet in a Local Zone, MGN launches the replication server in that Local Zone the same way it would in any Availability Zone.

When you replicate to a Local Zone, you can store Amazon EBS snapshots in the Local Zone instead of the parent AWS Region. By default, snapshots of Amazon EBS volumes in a Local Zone are stored in the parent AWS Region. If you replicate to a Local Zone that supports local snapshots, you can store the snapshots locally in the Local Zone to meet data residency requirements.

Local snapshots are supported only in Local Zones where MGN offers this feature. If MGN supports local snapshots in the Local Zone that contains your staging area subnet, the option to store snapshots locally is available. If MGN does not support the Local Zone, the option is not available and snapshots are stored in the parent AWS Region. Local snapshots are supported in the following Local Zones, and in additional Local Zones:

- Hanoi, Vietnam Local Zone (ap-southeast-1-han-1a), which has the Asia Pacific (Singapore) Region as its parent Region.
- Athens, Greece Local Zone (eu-central-1-ath-1a), which has the Europe (Frankfurt) Region as its parent Region.
- Istanbul, Türkiye Local Zone (eu-central-1-ist-1a), which has the Europe (Frankfurt) Region as its parent Region.

For more information about local snapshots in Local Zones, see [Local snapshots in Local Zones](#) in the Amazon EBS User Guide.

Always use AWS Transform MGN security group

Choose whether you would like to **Always use the AWS Transform MGN security group**.

A security group acts as a virtual firewall, which controls the inbound and outbound traffic of the staging area subnet.

The best practice is to have AWS Transform MGN automatically attach and monitor the default AWS Transform MGN security group. This group opens inbound TCP Port 1500 for receiving the transferred replicated data. When the default AWS Transform MGN security group is activated, MGN will constantly monitor whether the rules within this security group are enforced, in order to maintain uninterrupted data replication. If these rules are altered, MGN will automatically fix the issue.

Select the **Always use AWS Transform MGN security group** check box to allow data to flow from your source servers to the replication servers, and to allow the replication servers to communicate their state to the AWS Transform MGN servers.

To opt out, clear the check box. Clearing the check box is not recommended.

Additional security groups can be chosen from the Additional security groups dropdown. The list of available security groups changes according to the **Staging area subnet** you selected.

You can search for a specific security group within the search box.

You can add security groups via the AWS Management Console, and they will appear on the security group drop-down list in the AWS Transform MGN Console. Learn more about AWS security groups in [this VPC article](#).

You can use the default AWS Transform MGN security group, or you can select another security group. However, take into consideration that any selected security group that is not the AWS Transform MGN default, will be added to the default group, since the default security group is essential for the operation of MGN.

Data routing and throttling

AWS Transform MGN allows you to control how data is routed from your source servers to the replication servers on AWS through the **Data routing and throttling** settings.

By default, data is sent from the source servers to the replication servers over the public internet, using the public IP that was automatically assigned to the replication servers. Transferred data is always encrypted in transit.

 Note

The **Data routing and throttling** view differs slightly between the replication template view and the individual source server replication settings view, but the instructions apply to both views.

Use private IP for data replication

Choose the **Use private IP** option if you want to route the replicated data from your source servers to the staging area subnet through a private network with a VPN, AWS Direct Connect, VPC peering, or another type of existing private connection.

Choose **Do not use private IP** if you do not want to route the replicated data through a private network.

 Important

Data replication will not work unless you have already set up the VPN, AWS Direct Connect, or VPC peering in the AWS Console.

 Note

- Use of private IP is not supported for IPv6.
- If you selected the Default subnet, it is highly unlikely that the private IP is used for that subnet. Ensure that Private IP (VPN, AWS Direct Connect, or VPC peering) is used for your chosen subnet if you wish to use this option.
- You can safely switch between a private connection and a public connection for individual server settings choosing the **Use private IP** or **Do not use private IP** option, even after data replication has begun. This switch will only cause a short pause in replication, and will not have any long-term effect on the replication.
- Choosing the **Use private IP** option will not create a new private connection.


You should use this option if you want to:

- Allocate a dedicated bandwidth for replication

- Use another level of encryption
- Add another layer of security by transferring the replicated data from one private IP address (source) to another private IP address (on AWS)

Network architecture diagram – private IP

The following diagram illustrates the high-level interaction between the different replication system components when using private IP or VPC endpoint.



Create public IP

When the **Use private IP** option is chosen, you will have the option to create a public IP. Public IPs are used by default. Choose **Create public IP** if you want to create a public IP. Choose **Do not create a public IP** if you do not want to create a public IP.

Throttle bandwidth

You can control the amount of network bandwidth used for data replication per server. By default, AWS Transform MGN will use all available network bandwidth utilizing five concurrent connections.

Choose **Throttle bandwidth** if you want to control the transfer rate of data sent from your source servers to the Replication Servers over TCP Port 1500. Otherwise, choose **Do not throttle bandwidth**.

If you chose to throttle bandwidth, the **Throttle network bandwidth** (per server, in Mbps) box will appear. Enter your desired bandwidth in Mbps.

Replication resources tags

Add custom **Replication resources tags** to resources created by AWS Transform MGN in your AWS account.

These are resources required to facilitate data replication, testing and cutover. Each tag consists of a key and an optional value. You can add a custom tag to all of the AWS resources that are created on your AWS account during the normal operation of AWS Transform MGN.

To add a new tag, take the following steps:

1. Choose **Add new tag**.
2. Enter a **Custom tag key** and an optional tag value.

Note

MGN already adds tags to every resource it creates, including service tags and user tags.

These resources include:

- Amazon EC2 instances
- Amazon EC2 launch templates
- Amazon EBS volumes
- Snapshots
- Security groups (optional)

Learn more about AWS Tags in [this Amazon EC2 article](#).

Launch template

The **Launch template** allows you to control the way MGN launches instances in AWS. The default configuration defined in the template is automatically applied to every newly added server. You can modify the template at any time. Template settings are applied to each newly added server, but changes do not affect servers already in the MGN console. You can change the settings for existing and newly added servers individually within the server details view.

To edit the launch template for your entire account, you need to edit your launch template. Choose **Launch template** from the left-hand navigation menu.

This opens the account template view. Choose **Edit** to update your account-level launch template.

Note

- Ensure the Account Level EC2 Launch Template is not deleted. If it is, edit and save the launch template page to create a new Account Level Template.
- You can change the settings for specific servers without affecting the launch template as described in [Edit replication settings for a server](#).
- See [Assign an IPv6 address to an instance](#) in the *Amazon EC2 User Guide* for a list of instance types that do not support IPv6 addresses

Topics

- [General launch settings reference](#)
- [Default EC2 launch template settings](#)
- [MAP program tagging setting](#)

General launch settings reference

In the **General launch settings** section, you can define:

- **Instance type right sizing:** If you select this option, MGN launches a test or cutover AWS instance type that best matches the OS, CPU, and RAM of your source server. Note that the AWS instance type selected by MGN when this option is selected overwrites the instance type defined in your EC2 launch template.

- **Start instance upon launch:** Choose whether to start your test and cutover instances automatically upon launch or launch them in a stopped state.
- **Copy private IP:** Select if you want MGN to ensure that the private IP used by the test or cutover instance matches the private IP used by the source server. If selected, MGN monitors the source server on an hourly basis to identify the Private IP and uses the private IP of the primary network interface. Make sure that the IP range of the subnet you set in the EC2 Launch Template includes the private IP address for this feature to work. Copy private IP is not supported for IPv6.
- **Transfer server tags:** Select if you want MGN to transfer user-configured custom tags from your source servers to your test or cutover instance. If selected, server tags are transferred. These tags are attached to all source servers, all launched test and cutover instances, and all of the ephemeral resources that are created on your AWS account during the normal operation of MGN, such as snapshots, storage volumes (Amazon EBS and FSx for ONTAP), replication and conversion servers, and security groups.
- **Operating system licensing:** Select if you want to Bring Your Own Licenses (BYOL) from the source server into the test or cutover instance.

Select the “Bring your own license (**BYOL**)” option if you are migrating a Linux server. All Linux licenses are BYOL by default. Any RHEL, SUSE or Debian licenses are transferred in their current form to the migrated instance.

For Windows servers, choose whether you want to use **BYOL**. This sets up a Dedicated Host. All the licenses from the source Windows server are automatically transferred to the test or cutover instance. Note that if you use BYOL licensing for Windows you have to change the **Placement.tenancy** type in the EC2 launch template to **Host**. Otherwise, instance launch fails.

- **Boot mode:** When a computer boots, the first software that it runs is responsible for initializing the platform and providing an interface for the operating system to perform platform-specific operations. In Amazon EC2, two variants of the boot mode software are supported: Unified Extensible Firmware Interface (UEFI) and Legacy BIOS. The boot mode allows MGN to launch a test or cutover instance type that best matches the configuration of the source server. You can select between **Keep the source boot mode** or change it to **BIOS/UEFI modes**.

 Note

UEFI is not supported in CentOS 6 and RHEL 6.

Default EC2 launch template settings

In the **Default EC2 launch template** section, you can define the following options:

- **Default target subnet** – Choose the target subnet for the test and cutover machines. All the test and cutover machines are launched in this target subnet.
- **Target security groups** – Choose a number of security groups for the test and cutover machines. Upon the target launch, the selected security groups are attached to the Amazon EC2 instances.

Note

If you have selected FSx for ONTAP as the target storage in Replication settings, ensure you select a security group that allows your launched instances to connect to the FSx for ONTAP file system over iSCSI (TCP port 3260). See [Step 1: Configure security groups](#) in the [FSx for ONTAP configuration guide](#).

- **Default instance type** – Choose a default instance type for the test and cutover machines. The instance type chosen in this template is propagated to source server's launch template and is used to launch the target instance.
- **EBS volume type** – When using Amazon EBS as the target storage type, choose an Amazon EBS volume type for the test and cutover machines. The Amazon EBS volume type options are io1, io2, gp3, st1, and sc1. This setting does not apply when using FSx for ONTAP.

If this setting is not chosen, the default volume type selected in the launch template is gp3 with maximum IOPS (16000). If a volume type is selected in this option, and if the disk size does not match the limits of the volume type, the default gp3 volume type is used with maximum IOPS.

- **IOPS** – General Purpose SSD (gp3) volumes support a baseline of 3,000 IOPS. Additionally, you can provision up to 500 IOPS per GiB, up to a maximum of 16,000 IOPS. The default value is 3,000 IOPS.
- **Throughput** – General Purpose SSD (gp3) volumes have a baseline performance of 125 MiB/s. You can provision additional throughput of 0.25 MiB/s per provisioned IOPS, up to a maximum of 2,000 MiB/s (at 8,000 IOPS or higher).
- **Volume initialization rate** – Specify the volume initialization rate, in MiB/s, at which the snapshot data is downloaded to the volume. Specifying an initialization rate ensures that volumes are fully initialized in a predictable amount of time. This option is supported only for volumes created from snapshots, and additional costs apply. For more information, see [Amazon EBS volume initialization rate](#) in the *Amazon EBS User Guide*.

- **Delete on termination** – Indicates whether the volume should be automatically deleted when the instance is terminated. If you disable this feature, the volume persists independently from the running life of an Amazon EC2 instance. When you terminate the instance, the volume remains provisioned in your account. If you no longer need the volume after the instance has been terminated, you must delete it manually. You can also change the delete on termination behavior after the instance has been launched. For more information, see [How instance termination affects root and data volumes](#) in the *Amazon EC2 User Guide*.

Note

Note: if you manually delete the default launch template, MGN generates a new default launch template. Any changes previously made to the default template are discarded, including subnet and security groups. You can make the same changes on the new launch template, and they are applied to servers you add to the console after you make the changes.

MAP program tagging setting

Use this setting to determine whether to apply Migration Acceleration Program (MAP) tags to your launched instances. Learn more about the [AWS Migration Acceleration Program \(MAP\)](#).

Select **Add MAP tag to Launched Instances** option, if you want MGN to automatically tag your launched instances with the tag key and value combination required for MAP program. Once selected you must specify the MAP tag value that is used in your MAP tagging. MGN automatically tags your migrated resources with the key: "map-migrated" and the value of the tag that you provided. For more details about the tag value that should be used here, refer to the MAP tagging guide provided in your MAP term.

[Learn more about the AWS Migration Acceleration Program \(MAP\).](#)

Post-launch template

The Post-launch template allows you to control and automate actions performed after the server has been launched in AWS. The server settings are created automatically based on the **Post-launch template**.

To access the template, choose **Post-launch template** from the navigation menu.

The configuration defined in the template is applied to every newly added server. Changes to the template do not affect servers already available in the MGN console when the changes are made. You can change the settings for existing and newly added servers individually within the server details view.

Important

To use the post-launch settings feature, you must first activate at the account level. Deactivation is also at the account level.

The **Post-launch template** allows you to control various post-launch actions, including:

- Deployment of test and cutover instances
- Disaster recovery configuration (installing the AWS Replication Agent for Elastic Disaster Recovery and configuring the target disaster recovery AWS Region).
- Operating system conversion on the target machine
- License and subscription changes on the target machine

Topics

- [Activating post-launch settings](#)
- [Editing the post-launch settings template](#)
- [Deploying post-launch actions](#)
- [Encrypt post-launch action parameters](#)
- [Post-launch actions table](#)
- [Create a custom post-launch action](#)
- [Edit custom post-launch actions](#)
- [Predefined post-launch actions reference](#)

Activating post-launch settings

To use the post-launch template activate the post-launch actions. This allows MGN to:

- Install the [SSM Agent](#) on your servers

- Run the post-launch actions

Note

- Installation of the SSM Agent requires a minimum of 200 MB of free disk space and 200 KB of free disk space in the `/var` directory.
- Installation of the SSM Agent is not supported on these operating systems:
 - CentOS 5.x
 - CentOS 6.x
 - RHEL 6.x
 - Oracle 6.x
 - Amazon Linux 1
 - Windows 2008
 - Windows 2008 R2

To activate the post-launch actions:

1. Navigate to **Settings > Post-launch settings template**.
2. Choose **Edit**.
3. Toggle the **Install Systems Manager agent and allow executing actions on launched servers** option.
4. Choose **Save template**.

The post-launch actions are shown in the **Settings > Post-launch template** view.

Editing the post-launch settings template

MGN supports post-launch modernization actions, giving you the opportunity to move and improve. The service provides actions that you can execute on your Amazon EC2 launch instances and enables you to create your own actions.

The actions described in these sections can be edited within the post-launch template. Once you have edited your settings, choose **Save template**.

Deploying post-launch actions

Use this setting to choose whether to execute the post-launch actions on your cutover instances, on your test instances, or on both cutover and test instances.

Encrypt post-launch action parameters

The post-launch action parameters are stored in SSM [Parameter Store](#). For enhanced security, ensure that users who do not have permissions to execute SSM documents, do not have access to the Parameter Store. For an additional layer of security you can select to encrypt the action parameters using AWS KMS encryption.

SSM encrypts the parameter value of SecureString parameters type using AWS KMS with an AWS managed key or with the default AWS KMS key provided by AWS. You can specify different keys for each parameter, or use the same key for multiple parameters.

Post-launch actions table

The post-launch actions table includes both predefined actions and custom actions that are executed on your new Amazon EC2 instances.

- Predefined post-launch actions are provided out of the box. They are prepopulated with the necessary values and only need to be activated or deactivated. These actions are based on public SSM documents that cannot be changed and have certain unchangeable parameters such as the platform name and order.
- Custom post-launch actions are based on SSM documents that you create and upload to your account.

Use the **Filter by** options to filter the available actions according to your preferences.

Choose the settings icon to alternate between card and list view, according to your preferences.

Create a custom post-launch action

AWS Transform MGN allows you to execute any SSM document that you like – public SSM document or ones you created and uploaded to your account.

You can configure a custom action to execute any SSM document that is available in your account.

To add a new custom action, go to the **Post-launch actions settings** and choose **Create action**.

The page includes these parameters:

- **Action name** – The name of the action in MGN, which should be intuitive and meaningful to your migration users.
- **Activate this action** – Use this checkbox to activate or deactivate the custom action.
- **This action must be completed successfully before finalizing cutover** – This checkbox dictates whether or not the script prevents the cutover.
- **Systems Manager document name** – Select any SSM document that is available for the specific account.
- **View in Systems Manager** – Choose to open **SSM** and view additional information about the document.
- **Description** – Add a description or keep the default.
- **Document version** – Select which SSM document version to run. MGN can run a default version, the latest version, or a specific version, according to your preferences.
- **Category** – Select from various available categories including disaster recovery, security, validation, and more.
- **Order** – Specify the order in which the actions are executed. The lower the number, the earlier the action is executed. 1–1,000 are reserved for predefined actions and 1,001–10,000 for custom actions. The numbers must be unique but don't need to be consecutive.
- **Operating system** – Select the source server's operating systems for which the custom action can be configured for. Note that if you associate a script with the wrong operating system, it is skipped.
- **Creator** – Who created the action. For custom actions, the default is always **Me**.

The **Action parameters** change according to the specific SSM document that is selected.

Note that for the instance ID parameter, you can choose to use the launch instance ID, in which case, MGN dynamically populates the value.

Note

Only trusted, authorized users should have access to the parameter store. For enhanced security, ensure that users who do not have permissions to execute SSM documents / commands, do not have access to parameter store. [Learn more about restricting access to SSM parameters](#). Action parameters are stored in the SSM parameter store as regular

strings. Changing parameters in the SSM Parameter store may impact the post launch action execution on target instances. We recommend you consider security implications, when choosing to use parameters that contain scripts or sensitive information, such as API keys and database passwords.

Edit each setting as required and then choose **Add action**.

Edit custom post-launch actions

AWS Transform MGN allows you to execute any SSM document that you like – public SSM document or ones you created and uploaded to your account.

You can configure a custom action to execute any SSM document that is available in your account.

Use this page to edit the parameters detailed in the **Create action** section.

Edit each setting as required and then choose **Save action**.

Predefined post-launch actions reference

AWS Transform MGN allows you to execute various predefined post-launch actions on your Amazon EC2 launch instance. Use these out-of-the-box actions to modernize your servers while you're migrating: Change existing license, upgrade your operating system, configure disaster recovery, and more.

Choose from these predefined post-launch actions:

- [Install the SSM agent](#)
- [Configure AWS Elastic Disaster Recovery](#)
- [Convert operating systems](#)
- [Replace SUSE subscription](#)
- [Conduct Amazon EC2 connectivity checks](#)
- [Validate volume integrity](#)
- [Verify process status](#)
- [Convert MS-SQL license](#)
- [Install a CloudWatch Agent](#)
- [Upgrade Windows](#)

- [Create AMI from instance](#)
- [Join Directory Service domain](#)
- [Configure Time Sync](#)
- [Validate disk space](#)
- [Verify HTTP/HTTPS response](#)
- [Enable Amazon Inspector Classic](#)
- [Verify Tags](#)
- [Auto Scaling group setting](#)
- [Dynatrace](#)
- [New Relic](#)
- [TrendMicro](#)

Install the SSM agent

The SSM allows AWS Transform MGN to execute modernization actions on your servers after they are launched.

When you activate the post-launch actions, AWS Transform MGN installs the **SSM agent** and creates the required IAM roles.

The SSM agent must be installed for any other post-launch action to run. Therefore, this is the only post-launch action that is activated by default and cannot be deactivated.

[Learn more about SSM.](#)

Configure AWS Elastic Disaster Recovery

Note

This feature is supported on operating systems that are supported by AWS Elastic Disaster Recovery (AWS DRS). [See the AWS DRS documentation.](#)

This action is not supported in MGN GovCloud regions (US-East, US-West).

Use the **DR after migration** feature to configure disaster recovery using AWS Elastic Disaster Recovery.

This action installs the AWS Elastic Disaster Recovery Replication Agent on your Amazon EC2 instance.

You must select the target disaster recovery region, which is the AWS Region in which the Recovery instances are deployed. AWS Elastic Disaster Recovery must be available in the selected Region and initiated in your account. You must initialize Elastic Disaster Recovery for this action to work.

 Note

Ensure that you review the costs associated with AWS Elastic Disaster Recovery in the [service pricing documentation](#).

[Learn more about Elastic Disaster Recovery AWS Regions.](#)

[Learn more about initializing Elastic Disaster Recovery.](#)

Convert operating systems

 Note

This feature is supported on CentOS version 8.x.

Use the **CentOS to Rocky** feature to perform changes to the target machine operating system. It allows you to convert any of your source servers that are running CentOS to [Rocky Linux](#).

Replace SUSE subscription

 Note

- This feature is supported on SUSE Linux versions 12 SP 1 and later.
- This action is not supported on SLES4SAP servers.

Use the **Replace SUSE subscription** feature to choose whether you want to change the SUSE Linux subscription of any source server that runs SUSE to an AWS-provided SUSE subscription.

An AWS-provided SUSE subscription allows AWS to manage your licenses, including renewal handling, saving you time and simplifying your billing and license management processes

Conduct Amazon EC2 connectivity checks

Use the **EC2 connectivity check** feature to conduct network connectivity checks to a predefined list of ports and hosts.

Note

Up to 5 Port:IP couples can be checked in a single action.

Validate volume integrity

Use the **Volume integrity validation** feature to ensure that Amazon EBS volumes on the launched instance are:

- The same size as the source (rounded up)
- Properly mounted on the Amazon EC2 instance
- Accessible

For FSx for ONTAP migrations using iSCSI, this action additionally validates that:

- iSCSI sessions are established to the ONTAP SVM
- Multipath paths are active (when MPIO is configured)

This feature allows you to conduct the required validations automatically and saves the time of manual validations.

Note

Up to 50 volumes can be checked in a single action.

Verify process status

Use the **Process status validation** feature to ensure that processes are in running state following instance launch. You need to provide a list of processes that you want to verify, and define how long the service should wait before testing begins.

To check a specific process that should run multiple times, include it several times in the list.

Convert MS-SQL license

Use the **Windows MS-SQL license conversion** feature to easily convert Windows MS-SQL BYOL to an AWS license. MGN detects SQL Server installed as either the default instance (MSSQLSERVER) or as a named instance (for example, MSSQL\$SQLPROD).

MGN:

- Checks the SQL edition (Enterprise, Standard, or Web) as part of the launch process
- Uses the right AMI with the right billing code to launch from

The SSM document runs and verifies that the right billing code is used post launch.

The action uses these APIs:

- [DescribeImages](#)
- [DescribeInstances](#)

To allow the SSM document to run these APIs, you need the required permissions or have access to a role with those permissions and then provide the role's ARN as an input parameter to the SSM automation document.

Install a CloudWatch Agent

Use the **CloudWatch agent installation** feature to install and configure the CloudWatch Agent and Application Insights.

You need the `AWSApplicationMigrationSSMAccess` policy, or a user-defined policy that allows the SSM document to run, to run this post-launch action. This is in addition to the [full access policy](#):

The launched instance requires these policies:

- **CloudWatchAgentServerPolicy** – The permissions required to use AmazonCloudWatchAgent on servers
- **AmazonSSMManagedInstanceCore** – The policy for Amazon EC2 Role to enable AWS Systems Manager service core functionality

To ensure that the launch instance has the right policies, create a role that has the required permissions as per the policies above or has access to a role with those permissions.

- Go to **Launch settings > EC2 launch template > Modify > Advance > IAM instance profile**.
- Use an existing profile or create a new one using the **Create new IAM profile** link.

Note

- You must attach both policies to the template for the CloudWatch agent to operate. Without the CloudWatchAgentServerPolicy, the action is still marked as successful but the CloudWatch Agent is not active.
- Configuring the Application Insights is optional. You can choose to skip the Application Insights agent configuration and only install the CloudWatch agent. To do so provide the required parameterStoreName parameter and leave the other parameters empty.

[Learn more about the CloudWatch Agent.](#)

Upgrade Windows

Use the **Windows upgrade** feature to upgrade your migrated server to a more recent version of Windows Server ([see the full list of available OS versions](#)).

You need the AWSApplicationMigrationSSMAccess policy, or a user-defined policy that allows the SSM document to run, to run this post-launch action. This is in addition to the [full access policy](#):

To allow the SSM document to run these APIs, you must have the required permissions (including [CreateImage](#), [RunInstances](#), [DescribeInstances](#), and more) or have access to a role with those permissions and then provide the role's ARN as an input parameter to the SSM automation document.

Learn more about the permissions required to perform the upgrade in [AWSEC2-CloneInstanceAndUpgradeWindows](#).

The SSM document:

- Creates an Amazon Machine Image (AMI) from the instance using the [CreateImage](#) API.
- Uses the AMI to create a new instance and then upgrades that instance.
- Creates an AMI from the upgraded instance and terminates the upgraded instance.

 Note

- This operation may run for several hours.
- All other post-launch actions run on the instance launched by MGN and not on the upgraded instance.

[Learn more about upgrading Windows.](#)

Create AMI from instance

Use the **Create AMI from Instance** feature to create a new Amazon Machine Image (AMI) from your MGN launched instance.

You need the `AWSApplicationMigrationSSMAccess` policy, or a user-defined policy that allows the SSM document to run, to run this post-launch action. This is in addition to the [full access policy](#):

The action uses these APIs:

- [CreateImage](#)
- [DescribeImages](#)

To allow the SSM document to run these APIs, you need the required permissions or have access to a role with those permissions and then provide the role's ARN as an input parameter to the SSM automation document.

[Learn more about creating AMI from instance.](#)

Join Directory Service domain

Use this **Join domain** feature to simplify the AWS Join Domain process. If you activate this action, your instance is managed by the AWS Cloud Directory (instead of on-premises).

You need the `AWSApplicationMigrationSSMAccess` policy, or a user-defined policy that allows the SSM document to run, to run this post-launch action. This is in addition to the [full access policy](#):

The launched instance requires these policies:

- `AmazonSSMManagedInstanceCore` – The policy for Amazon EC2 Role to enable AWS Systems Manager service core functionality.
- `AmazonSSMDirectoryServiceAccess` – This policy allows the SSM Agent to access Directory Service on behalf of the customer for domain-join the managed instance.

To ensure that the launched instance has the right policies, create a role that has the required permissions as per the policies above or has access to a role with those permissions.

- Go to **Launch settings > EC2 launch template > Modify > Advance > IAM instance profile**.
- Use an existing profile or create a new one using the **Create new IAM profile** link.

Configure Time Sync

Use the **Time Sync** feature to set the time for your Linux instance using ATSS.

[Learn more about Amazon Time Sync.](#)

Validate disk space

Use the **Disk space validation** feature to obtain visibility into the disc space that you have at your disposal, as well as logs with actionable insights.

Verify HTTP/HTTPS response

Use the **Verify HTTP/HTTPS response** feature to conduct HTTP/HTTPS connectivity checks to a predefined list of URLs. The feature verifies that HTTP/HTTPS requests (for example, `https://localhost`) receive the correct response.

Enable Amazon Inspector Classic

The **Enable Inspector** feature allows you to run security scans on your Amazon EC2 resources. The Amazon Inspector service is enabled at the account level.

Note

Amazon Inspector is a paid AWS service. For additional information, [refer to the full Inspector pricing documentation](#).

This action uses these APIs:

- [Enable](#)
- [BatchGetAccountStatus](#)
- [CreateServiceLinkedRole](#)

To allow the SSM document to run these APIs, you need the required permissions or have access to a role with those permissions and then provide the role's ARN as an input parameter to the SSM automation document.

Verify Tags

Use the **Verify tags** feature to validate that tags that have been defined in the launch template and on the source server are copied to the migrated server.

Auto Scaling group setting

Use the **Auto Scaling group setting** when you would like to create an Auto Scaling group for a migrated stateless web application.

Dynatrace

Note

This action is provided by a third party vendor, and is not available in the GovCloud Regions.

This action installs Dynatrace OneAgent on your launched instance.

To configure this action, you need an existing Dynatrace account and configure the required `additionalArguments` for your particular usage.

Learn more about Dynatrace in [Deploy OneAgent using AWS Systems Manager Distributor](#)

New Relic

Note

This action is provided by a third party vendor, and is not available in the GovCloud Regions.

This action installs New Relic Infrastructure agent on your launched Amazon EC2 instance.

To configure this action, you need an existing New Relic account and configure the required `additionalArguments` for your particular usage. You must use an original account license key for this action to succeed.

[Learn more about New Relic](#)

TrendMicro

Note

This action is provided by a third party vendor, and is not available in the GovCloud Regions.

This action installs the Trend Micro agent on your launched instance.

[Learn more about Trend Micro](#)

Network requirements for MGN

Before you use MGN, make sure to prepare your environments. Preparation includes setting correct network settings, defining network requirements, and opening the correct ports.


Topics

- [Service and network architecture overview](#)
- [Network setting preparations](#)
- [Required connectivity settings](#)
- [Troubleshooting connectivity issues](#)

Service and network architecture overview

Watch the [AWS Transform MGN - Service architecture and network architecture video](#) for an in-depth overview of the MGN architecture.

This is the MGN network diagram:



Network setting preparations

As part of network setting preparations, set up a staging area and operational subnets.

Topics

- [Staging area subnet](#)
- [Operational subnets](#)
- [Active Directory connectivity](#)

Staging area subnet

Before setting up MGN you should create a subnet which will be used by the service as a staging area for data replicated from your source servers to AWS.

- You must specify this subnet in the replication settings template. While you can use an existing subnet in your AWS account, the best practice is to create a new dedicated subnet for this purpose. [Learn more about replication settings.](#)
- You can override this subnet for specific source servers in the [replication settings](#).

Note

MGN does not require special configuration for AWS Local Zones. You specify a staging area subnet and a launch subnet. MGN uses whichever subnet you provide, regardless of whether it is in an Availability Zone or a Local Zone.

For optimal performance of replication servers and conversion servers, we recommend placing the staging area subnet in the parent AWS Region rather than the Local Zone. If your requirements dictate using a Local Zone subnet for the staging area, we recommend testing replication and cutover thoroughly before proceeding with production migrations.

Operational subnets

Test and cutover instances are launched in a subnet you specify in the Amazon EC2 launch template associated with each source server. The Amazon EC2 launch template is created automatically when you add a source server to AWS Transform MGN.

[Learn more about launching test and cutover instances.](#)

[Learn more about how Amazon EC2 launch templates are used.](#)

Active Directory connectivity

If your source servers are domain-joined to an Active Directory, ensure that the target VPC has network connectivity and DNS resolution to your AD domain controllers before launching test or cutover instances. During conversion, MGN resets network settings to use DHCP, which means instance-level DNS overrides from the source server are not preserved.

For details on configuring DNS and network connectivity for domain-joined servers, see [Target instance cannot connect to Active Directory after migration.](#)

⚠ Important

Isolate your test VPC from your production network. If you launch a domain controller into a test VPC that has a network route back to your on-premises or production environment, production clients might bind to the launched domain controller instead of your production Active Directory, which can disrupt production workloads. When testing, we recommend that you do not launch Active Directory domain controllers into an environment that has connectivity to production, and that you remove any routes between your test VPC and your production network.

Cutover is different. A cutover instance is intended to replace your production server, so connectivity to your existing environment is expected. When you cut over a domain controller, plan the timing carefully: coordinate the shutdown of the corresponding on-premises or source domain controller with the cutover so that clients transition to the migrated domain controller in a controlled way, rather than having two active domain controllers serve the same clients at once.

If you plan to operate in a hybrid configuration, where your on-premises environment remains online alongside servers migrated to the cloud, plan network connectivity between the two environments carefully. Ensure that clients resolve and connect to the intended Active Directory domain controllers so that you avoid unexpected results, such as clients binding to the wrong domain controller.

Required connectivity settings

To prepare your network for running AWS Transform MGN, configure the following connectivity settings. All communication is encrypted with TLS.

⚠ Important

SSL interception should not be applied for communication between replication servers and the MGN API endpoint, or between source servers and the MGN API endpoint.

Topics

- [Endpoints and firewall allowlist](#)
- [Communication between source servers and AWS Transform MGN over TCP port 443](#)

- [Communication between source servers and the staging area subnet over TCP port 1500](#)
- [Communication between the staging area subnet and AWS Transform MGN over TCP port 443](#)
- [Communication between the staging area subnet and Amazon S3](#)
- [Network requirements for FSx for ONTAP](#)

Endpoints and firewall allowlist

Add the following endpoints and URLs to your firewall allowlist.

MGN API endpoints

Protocol	Endpoint
IPv4	<code>https://mgn. <i>region</i>.amazonaws.com</code>
Dual-stack (IPv4/IPv6)	<code>https://mgn. <i>region</i>.api.aws</code>

Replace *region* with the AWS Region code you are replicating to (for example, us-east-1).

Amazon EC2 endpoint

The staging area subnet requires outbound access to the [Amazon EC2 endpoint of its AWS Region](#).

Amazon S3 buckets

The AWS Replication Agent installer and the staging area subnet require access to Amazon S3. Allowlist these buckets:

```
https://aws-mgn-internal-<region>.s3.<region>.amazonaws.com/
https://aws-mgn-internal-hashes-<region>.s3.<region>.amazonaws.com/
https://aws-application-migration-service-<region>.s3.<region>.amazonaws.com/
https://aws-application-migration-service-hashes-<region>.s3.<region>.amazonaws.com/
https://amazon-ssm-<region>.s3.<region>.amazonaws.com/
https://al2023-repos-<region>-de612dc2.s3.dualstack.<region>.amazonaws.com/
```

Note

For IPv6/dual-stack environments, use the Amazon S3 dual-stack endpoint format: `https://<bucket>.s3.dualstack.<region>.amazonaws.com/`. Learn more about [Amazon S3 dual-stack endpoints](#).

Important

Ensure that your Amazon S3 VPC gateway endpoint policy allows access to the Amazon Linux 2023 package repository bucket (a12023-repos-<region>-de612dc2). This bucket is accessed through the Amazon S3 dual-stack endpoint.

Amazon S3 VPC Endpoint policy

If you use an Amazon S3 VPC Endpoint, you must provide sufficient permissions for service functionality, as shown in this example policy for replicating to us-east-1:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "*"
      },
      "Action": "s3:GetObject",
      "Resource": [
        "arn:aws:s3:::aws-mgn-internal-us-east-1/*",
        "arn:aws:s3:::aws-mgn-internal-hashes-us-east-1/*",
        "arn:aws:s3:::aws-application-migration-service-us-east-1/*",
        "arn:aws:s3:::aws-application-migration-service-hashes-us-east-1/*",
        "arn:aws:s3:::amazon-ssm-us-east-1/*",
        "arn:aws:s3:::a12023-repos-us-east-1-de612dc2/*"
      ]
    }
  ]
}
```

```
}
```

Communication between source servers and AWS Transform MGN over TCP port 443

Each source server that is added to MGN must continuously communicate with the MGN API endpoint over TCP port 443.

The main operations performed through this route are:

- Downloading and upgrading the AWS Replication Agent.
- Connecting the source servers to the MGN console and displaying their replication status.
- Monitoring the source servers for troubleshooting and resource consumption metrics (CPU, RAM).
- Reporting source server-related events (for example, disk changes).
- Transmitting source server information to the MGN console (hardware information, running services, installed applications and packages).
- Preparing the source servers for test or cutover.

Communication between source servers and the staging area subnet over TCP port 1500

Each source server with an installed AWS Replication Agent continuously communicates with the replication servers in the staging area subnet over TCP port 1500. This port is used for the transfer of replicated data from the source servers to the staging area subnet.

The replicated data is encrypted and compressed when transferred over TCP port 1500. The data is encrypted on the source infrastructure before transfer and decrypted at the staging area subnet before being written to the volumes.

AWS Transform MGN uses TLS 1.2 end to end from the agent installed on the source server to the Replication Server. Each replication server gets assigned a specific TLS server certificate, which is distributed to the corresponding Agent and validated on the agent side.

Communication between the staging area subnet and AWS Transform MGN over TCP port 443

The replication servers in the staging area subnet must continuously communicate with MGN over TCP port 443. The main operations performed through this route are:

- Downloading the replication software.
- Connecting the replication servers to MGN and displaying their replication status.
- Monitoring the replication servers for troubleshooting and resource consumption metrics.
- Reporting replication-related events.

There are two ways to establish direct connectivity to the Internet for the VPC of the staging area, as described in the [VPC FAQ](#):

1. [Public IP address + Internet gateway](#)
2. [Private IP address + NAT gateway](#)

Communication between the staging area subnet and Amazon S3

The staging area subnet requires access to Amazon S3 for both replication and conversion operations. Replication servers and conversion servers are based on Amazon Linux 2023 (AL2023) and require access to the AL2023 package repository hosted in Amazon S3 for package installation and updates.

The AL2023 package repository bucket is accessed through the Amazon S3 dual-stack endpoint. You must ensure that the staging area subnet can reach the following Amazon S3 URL:

```
https://al2023-repos-<region>-de612dc2.s3.dualstack.<region>.amazonaws.com/
```

Replace <region> with the AWS Region code you are replicating to (for example, us-east-1).

Depending on your network configuration, ensure access using one or both of the following methods:

- **Firewall-restricted environments** – If the staging area subnet has outbound internet access through a firewall or proxy, add the AL2023 repository URL to your allowlist. For the complete list of Amazon S3 URLs to allowlist, see [Endpoints and firewall allowlist](#).

- **Isolated subnets (no internet access)** – If the staging area subnet does not have outbound internet access and uses an Amazon S3 VPC gateway endpoint, ensure that the endpoint policy allows access to the AL2023 package repository bucket. Add the following resource ARN to the endpoint policy:

```
"arn:aws:s3:::al2023-repos-<region>-de612dc2/*"
```

For the complete Amazon S3 VPC endpoint policy example, see [Endpoints and firewall allowlist](#).

Important

Without access to the AL2023 package repository, replication servers may fail to launch or function correctly, and conversion servers may fail during the boot conversion process.

Network requirements for FSx for ONTAP

When using FSx for ONTAP as the target storage type, the following additional network connectivity is required between the target instances and the FSx for ONTAP file system.

Port	Protocol	Direction	Purpose
3260	TCP	Target instance → FSx for ONTAP	iSCSI data transfer for block storage replication
443	TCP	Target instance → FSx for ONTAP	ONTAP REST API for storage management (certificate-based authentication)

These ports must be allowed in the security groups attached to both the target instances and the FSx for ONTAP file system. For detailed security group configuration, see [Step 1: Configure security groups](#) in the [FSx for ONTAP configuration](#).

Replication server package repository access

MGN replication servers require iSCSI initiator and multipath packages to replicate data to FSx for ONTAP. These packages are installed automatically using yum from the Amazon Linux 2023 (AL2023) package repository.

- **Connected subnets (internet via NAT or internet gateway):** Allowlist the following URLs in your firewall or DNS rules:

```
https://cdn.amazonlinux.com/
https://al2023-repos-<region>-de612dc2.s3.dualstack.<region>.amazonaws.com/
```

- **Isolated subnets (Amazon S3 VPC gateway endpoint only):** Add the following resource ARN to the endpoint policy:

```
"arn:aws:s3:::al2023-repos-<region>-de612dc2/*"
```

If neither option is available, you must pre-install the packages on the source server before migration. For the required packages by operating system, see [Step 6: Configure launch template and launch settings](#).

Target instance package repository access

MGN automatically installs iSCSI initiator and multipath packages on test and cutover instances using the OS package manager of the target operating system. The launch subnet must have outbound access to the appropriate OS package repositories.

Required packages by package manager (Linux)

Package Manager	Packages Installed
dnf (Fedora/RHEL 8+)	iscsi-initiator-utils , device-mapper-multipath
yum (RHEL 6/7, CentOS, Amazon Linux)	iscsi-initiator-utils , device-mapper-multipath
apt-get (Debian/Ubuntu)	open-iscsi , multipath-tools
zypper (SLES/openSUSE)	open-iscsi , multipath-tools

On Windows, the iSCSI initiator (MSiSCSI service) is a built-in service that is enabled and started automatically. Only Multipath-IO needs to be enabled:

Required features (Windows)

Method	Feature Enabled
Install-WindowsFeature (Server 2012+)	Multipath-IO
Add-WindowsFeature (Server 2008 R2)	Multipath-IO

If the launch subnet cannot reach OS package repositories (for example, air-gapped environments or private subnets without a NAT gateway), or if the operating system uses subscription-based repositories (SUSE, RHEL, CentOS), you must pre-install the packages on the source server before migration. See the [Supported Linux operating systems](#) table for the specific pre-install commands by operating system.

Troubleshooting connectivity issues

If there is no connection between your source servers and MGN, make sure that your corporate firewall enables connectivity from the source servers to MGN over TCP port 443. If the connectivity is blocked, enable it.

Enabling Windows Firewall for TCP port 443 connectivity

Important

The information provided in this section is for general security and firewall guidance only. The information is provided on "AS IS" basis, with no guarantee of completeness, accuracy or timeliness, and without warranty or representations of any kind, expressed or implied. In no event will AWS and/or its subsidiaries and/or their employees or service providers be liable to you or anyone else for any decision made or action taken in reliance on the information provided here or for any direct, indirect, consequential, special or similar damages (including any kind of loss), even if advised of the possibility of such damages.

AWS is not responsible for the update, validation, or support of security and firewall information.

Note

These instructions are intended for the default OS firewall and allow outbound connectivity. You may still need to adjust other external components, such as firewall blocking or incorrect routes. Consult the documentation of any third-party local firewall you use.

1. On the source server, open the **Windows Firewall** console.
2. On the console, select the **Outbound Rules** option from the tree.

Outbound Rules											
Name	Group	Profile	Enabled	Action	Override	Program	Local Address	Remote Address	Protocol	Local Port	Remote Port
✓ AllJoyn Router (TCP-Out)	AllJoyn Router	Domai...	Yes	Allow	No	%System...	Any	Any	TCP	Any	Any
✓ AllJoyn Router (UDP-Out)	AllJoyn Router	Domai...	Yes	Allow	No	%System...	Any	Any	UDP	Any	Any
BranchCache Content Retrieval (HTTP-O...	BranchCache - Content Retr...	All	No	Allow	No	SYSTEM	Any	Any	TCP	Any	80
BranchCache Hosted Cache Client (HTTP...	BranchCache - Hosted Cach...	All	No	Allow	No	SYSTEM	Any	Any	TCP	Any	80, 443
BranchCache Hosted Cache Server(HTTP...	BranchCache - Hosted Cach...	All	No	Allow	No	SYSTEM	Any	Any	TCP	80, 443	Any
BranchCache Peer Discovery (WSD-Out)	BranchCache - Peer Discove...	All	No	Allow	No	%system...	Any	Local subnet	UDP	Any	3702
✓ Captive Portal Flow	Captive Portal Flow	All	Yes	Allow	No	Any	Any	Any	Any	Any	Any
✓ Cast to Device functionality (qWave-TCP...	Cast to Device functionality	Private...	Yes	Allow	No	%System...	Any	PlayTo Renderers	TCP	Any	2177
✓ Cast to Device functionality (qWave-UDP...	Cast to Device functionality	Private...	Yes	Allow	No	%System...	Any	PlayTo Renderers	UDP	Any	2177
✓ Cast to Device streaming server (RTP-Stre...	Cast to Device functionality	Private	Yes	Allow	No	%System...	Any	Local subnet	UDP	Any	Any
✓ Cast to Device streaming server (RTP-Stre...	Cast to Device functionality	Public	Yes	Allow	No	%System...	Any	PlayTo Renderers	UDP	Any	Any
✓ Cast to Device streaming server (RTP-Stre...	Cast to Device functionality	Domain	Yes	Allow	No	%System...	Any	Any	UDP	Any	Any
✓ Core Networking - DNS (UDP-Out)	Core Networking	All	Yes	Allow	No	%System...	Any	Any	UDP	Any	53
✓ Core Networking - Dynamic Host Config...	Core Networking	All	Yes	Allow	No	%System...	Any	Any	UDP	68	67
✓ Core Networking - Dynamic Host Config...	Core Networking	All	Yes	Allow	No	%System...	Any	Any	UDP	546	547
✓ Core Networking - Group Policy (LSASS-...	Core Networking	Domain	Yes	Allow	No	%System...	Any	Any	TCP	Any	Any
✓ Core Networking - Group Policy (NP-Out)	Core Networking	Domain	Yes	Allow	No	System	Any	Any	TCP	Any	445
✓ Core Networking - Group Policy (TCP-Out)	Core Networking	Domain	Yes	Allow	No	%System...	Any	Any	TCP	Any	Any
✓ Core Networking - Internet Group Mana...	Core Networking	All	Yes	Allow	No	System	Any	Any	IGMP	Any	Any
✓ Core Networking - IPHTTPS (TCP-Out)	Core Networking	All	Yes	Allow	No	%System...	Any	Any	TCP	Any	IPHTTPS
✓ Core Networking - IPv6 (IPv6-Out)	Core Networking	All	Yes	Allow	No	System	Any	Any	IPv6	Any	Any

3. On the **Outbound Rules** table, select the rule that relates to the connectivity to Remote Port - 443. Check if the **Enabled** status is **Yes**.

Outbound Rules											
Name	Group	Profile	Enabled	Action	Override	Program	Local Address	Remote Address	Protocol	Local Port	Remote Port
✓ AllJoyn Router (TCP-Out)	AllJoyn Router	Domai...	Yes	Allow	No	%System...	Any	Any	TCP	Any	Any
✓ AllJoyn Router (UDP-Out)	AllJoyn Router	Domai...	Yes	Allow	No	%System...	Any	Any	UDP	Any	Any
BranchCache Content Retrieval (HTTP-O...	BranchCache - Content Retr...	All	No	Allow	No	SYSTEM	Any	Any	TCP	Any	80
BranchCache Hosted Cache Client (HTTP...	BranchCache - Hosted Cach...	All	No	Allow	No	SYSTEM	Any	Any	TCP	Any	80, 443
BranchCache Hosted Cache Server(HTTP...	BranchCache - Hosted Cach...	All	No	Allow	No	SYSTEM	Any	Any	TCP	80, 443	Any
BranchCache Peer Discovery (WSD-Out)	BranchCache - Peer Discove...	All	No	Allow	No	%system...	Any	Local subnet	UDP	Any	3702
✓ Captive Portal Flow	Captive Portal Flow	All	Yes	Allow	No	Any	Any	Any	Any	Any	Any
✓ Cast to Device functionality (qWave-TCP...	Cast to Device functionality	Private...	Yes	Allow	No	%System...	Any	PlayTo Renderers	TCP	Any	2177
✓ Cast to Device functionality (qWave-UDP...	Cast to Device functionality	Private...	Yes	Allow	No	%System...	Any	PlayTo Renderers	UDP	Any	2177
✓ Cast to Device streaming server (RTP-Stre...	Cast to Device functionality	Private	Yes	Allow	No	%System...	Any	Local subnet	UDP	Any	Any
✓ Cast to Device streaming server (RTP-Stre...	Cast to Device functionality	Public	Yes	Allow	No	%System...	Any	PlayTo Renderers	UDP	Any	Any
✓ Cast to Device streaming server (RTP-Stre...	Cast to Device functionality	Domain	Yes	Allow	No	%System...	Any	Any	UDP	Any	Any
✓ Core Networking - DNS (UDP-Out)	Core Networking	All	Yes	Allow	No	%System...	Any	Any	UDP	Any	53
✓ Core Networking - Dynamic Host Config...	Core Networking	All	Yes	Allow	No	%System...	Any	Any	UDP	68	67
✓ Core Networking - Dynamic Host Config...	Core Networking	All	Yes	Allow	No	%System...	Any	Any	UDP	546	547
✓ Core Networking - Group Policy (LSASS-...	Core Networking	Domain	Yes	Allow	No	%System...	Any	Any	TCP	Any	Any
✓ Core Networking - Group Policy (NP-Out)	Core Networking	Domain	Yes	Allow	No	System	Any	Any	TCP	Any	445
✓ Core Networking - Group Policy (TCP-Out)	Core Networking	Domain	Yes	Allow	No	%System...	Any	Any	TCP	Any	Any
✓ Core Networking - Internet Group Mana...	Core Networking	All	Yes	Allow	No	System	Any	Any	IGMP	Any	Any
✓ Core Networking - IPHTTPS (TCP-Out)	Core Networking	All	Yes	Allow	No	%System...	Any	Any	TCP	Any	IPHTTPS
✓ Core Networking - IPv6 (IPv6-Out)	Core Networking	All	Yes	Allow	No	System	Any	Any	IPv6	Any	Any

- If the Enabled status of the rule is **No**, open the context menu for it and select **Enable Rule** from the pop-up menu.

Outbound Rules											
Name	Group	Profile	Enabled	Action	Override	Program	Local Address	Remote Address	Protocol	Local Port	Remote Port
✓ AllJoyn Router (TCP-Out)	AllJoyn Router	Domai...	Yes	Allow	No	%System...	Any	Any	TCP	Any	Any
✓ AllJoyn Router (UDP-Out)	AllJoyn Router	Domai...	Yes	Allow	No	%System...	Any	Any	UDP	Any	Any
BranchCache Content Retrieval (HTTP-O...	BranchCache - Content Retr...	All	No	Allow	No	SYSTEM	Any	Any	TCP	Any	80
BranchCache Hosted Cache Client (HTTP...	BranchCache - Hosted Cach...	All	No	Allow	No			Any	TCP	Any	80, 443
BranchCache Hosted Cache Server(HTTP...	BranchCache - Hosted Cach...	All	No	Allow	No			Any	TCP	80, 443	Any
BranchCache Peer Discovery (WSD-Out)	BranchCache - Peer Discove...	All	No	Allow	No			Local subnet	UDP	Any	3702
✓ Captive Portal Flow	Captive Portal Flow	All	Yes	Allow	No			Any	Any	Any	Any
✓ Cast to Device functionality (qWave-TCP...	Cast to Device functionality	Private...	Yes	Allow	No			PlayTo Renderers	TCP	Any	2177
✓ Cast to Device functionality (qWave-UDP...	Cast to Device functionality	Private...	Yes	Allow	No			PlayTo Renderers	UDP	Any	2177
✓ Cast to Device streaming server (RTP-Stre...	Cast to Device functionality	Private	Yes	Allow	No			Local subnet	UDP	Any	Any
✓ Cast to Device streaming server (RTP-Stre...	Cast to Device functionality	Public	Yes	Allow	No			PlayTo Renderers	UDP	Any	Any
✓ Cast to Device streaming server (RTP-Stre...	Cast to Device functionality	Domain	Yes	Allow	No			Any	UDP	Any	Any
✓ Core Networking - DNS (UDP-Out)	Core Networking	All	Yes	Allow	No	%System...	Any	Any	UDP	Any	53
✓ Core Networking - Dynamic Host Config...	Core Networking	All	Yes	Allow	No	%System...	Any	Any	UDP	68	67
✓ Core Networking - Dynamic Host Config...	Core Networking	All	Yes	Allow	No	%System...	Any	Any	UDP	546	547
✓ Core Networking - Group Policy (LSASS-...	Core Networking	Domain	Yes	Allow	No	%System...	Any	Any	TCP	Any	Any
✓ Core Networking - Group Policy (NP-Out)	Core Networking	Domain	Yes	Allow	No	System	Any	Any	TCP	Any	445
✓ Core Networking - Group Policy (TCP-Out)	Core Networking	Domain	Yes	Allow	No	%System...	Any	Any	TCP	Any	Any
✓ Core Networking - Internet Group Mana...	Core Networking	All	Yes	Allow	No	System	Any	Any	IGMP	Any	Any
✓ Core Networking - IPHTTPS (TCP-Out)	Core Networking	All	Yes	Allow	No	%System...	Any	Any	TCP	Any	IPHTTPS
✓ Core Networking - IPv6 (IPv6-Out)	Core Networking	All	Yes	Allow	No	System	Any	Any	IPv6	Any	Any

Enabling Linux Firewall for TCP port 443 connectivity

- Enter this command to add the required Firewall rule:

```
sudo iptables -A OUTPUT -p tcp --dport 443 -j ACCEPT
```

- To verify the creation of the Firewall rule, enter this command:

```
sudo iptables -L
```

Expected output:

```
Chain INPUT (policy ACCEPT)
target     prot opt source      destination
Chain FORWARD (policy ACCEPT)
target     prot opt source      destination
Chain OUTPUT (policy ACCEPT)
target     prot opt source      destination
ACCEPT     tcp  --  anywhere    anywhere      tcp dpt:443
```

Source Servers in AWS Transform MGN

You must add your source servers to the AWS Transform MGN console in order to migrate them into AWS. Source servers are added by installing the AWS Replication Agent on each individual server. This documentation provides installation paths for both Linux and Windows servers. Ensure that your servers are supported by AWS Transform MGN by reviewing the [Supported operating systems](#)

In addition, AWS Transform MGN allows you to perform [agentless snapshot replication](#) from your vCenter source environment into AWS. This is achieved by installing the MGN vCenter Client in your vCenter environment. Agentless snapshot replication has its own unique credentials, installation, and replication mechanisms. AWS Transform MGN recommends using agent-based replication when possible, as it supports CDP (Continuous Data Protection) and provides the shortest cutover window. Agentless replication should be used when your company's policies or technical issues prevent you from installing the AWS Replication Agent on each individual server.

Once your source servers have been added to AWS Transform MGN, you can monitor and interact with them from the Source servers page. The Source servers page is the default view in the AWS Transform MGN console, and is the page that you interact with the most. On the Source servers page, you can view all of your source servers, monitor their migration lifecycle and data replication state, see the next step in the migration process for each server, and sort your servers by a variety of categories. You can also perform a variety of commands from the Source servers page through the command menus. These menus allow you to fully control your servers by managing data replication, launching test and cutover instances, and disconnecting servers from AWS Transform MGN.

You can choose any individual source server on the Source servers page in order to access the server details view. This view allows you to see the details for individual servers. Here you are able to see an in-depth view of the server's migration lifecycle, browse an overview of the server's technical details, manage tags, manage disks, and most importantly, configure the individual replication settings and launch settings for the server.

Topics

- [Adding source servers](#)
- [Installing the AWS Transform MGN vCenter Client for Agentless Replication on vCenter source environments](#)
- [Manage source servers](#)

- [Access details on a source server](#)

Adding source servers

Add source servers to AWS Transform MGN by installing the AWS Replication Agent (also referred to as "the Agent") on them. The Agent can be installed on both Linux and Windows servers. You can add source servers from vCenter without installing an agent through the agentless replication feature.

Quick links:

- [Linux installation instructions](#)
- [Windows installation instructions](#)
- [Agentless replication instructions](#)

Note

While the use of AWS Transform MGN is free for 90 days, you will incur charges for any AWS infrastructure that is provisioned during migration and after cutover, such as compute (Amazon EC2) and storage (Amazon EBS or FSx for ONTAP) resources. These are billed to your account separately, at your regular rates.

Topics

- [Installation requirements](#)
- [Operating systems supported by MGN](#)
- [Installing the AWS Replication Agent](#)

Installation requirements

Before installing the AWS Replication Agent on your source servers, ensure that they meet these requirements:

General requirements

- Ensure that the source server operating system is supported by AWS.

- [Supported Windows operating systems.](#)
- [Supported Linux operating systems.](#)
- Ensure that your setup meets all networking requirements. [Learn more about network requirements.](#)
- Ensure MAC address stability – ensure that the MAC addresses of the source servers do not change upon a reboot or any other common changes in your network environment. AWS Transform MGN calculates the unique ID of the source server from the MAC address. When a MAC address changes, MGN is no longer able to correctly identify the source server. Consequently, replication stops. If this happens, you need to reinstall the AWS Replication Agent and start replication from the beginning.
- AWS Transform MGN does not support fully paravirtualized source servers. Source servers with partial paravirtualization, such as VMWare's paravirtualization of I/O devices, are supported.
- The AWS Replication Agent installer supports multipath.

Source server requirements

These are universal requirements for both Linux and Windows source servers:

- Root directory – Verify that your source server has at least 2 GB of free disk space on the root directory (/) .
- Verify that your source server has at least 300 MB of free RAM to run the AWS Replication Agent.
- MGN only supports operating systems built for the x86 system architecture.
- MGN supports replication of volumes up to a maximum size of 16 TB. MGN does not currently support migration of servers that have larger volume sizes. Agent installation will fail with an appropriate error message if the installer detects a volume that is larger than this limit.

Linux installation requirements

Ensure that your Linux source server meets these installation requirements prior to installing the AWS Replication Agent:

- Python is installed on the server – Python 2 (2.4 or above) or Python 3 (3.0 or above).
- These tools are required for agent installation only. The installer attempts to install them if they are not present already:

```
make gcc perl tar gawk rpm
```

- Verify that you meet these disk space requirements:
 - At least 2 GB of free disk space on the root directory (/) of your source server for the installation. To check the available disk space on the root directory, run the `df -h /` command.
 - At least 1 GB of free disk space on the `/tmp` directory for the duration of the installation process. To check the available disk space on the `/tmp` directory run the `df -h /tmp` command.
 - If `/boot` is a separate partition, ensure that it has a minimum of 50 MB free space needed for the installation. To check the available disk space on the `/boot` directory run the `df -h /boot` command.

After you have entered the commands for checking the available disk space, the results are displayed as:

```
ubuntu@Linux-1:~$ df -h /
Filesystem      Size  Used Avail Use% Mounted on
/dev/xvda1      7.8G  1.4G  6.0G  19% /
ubuntu@Linux-1:~$ df -h /tmp
Filesystem      Size  Used Avail Use% Mounted on
/dev/xvda1      7.8G  1.4G  6.0G  19% /tmp
```

- Ensure that you have Python installed on the source server (version 2.4+, version 3.0+) for Agent installation.
- Only servers using the GRUB bootloader (GRUB 1 or 2) are supported.
- Machines that boot off a disk configured with GPT partitioning must have the package 'grub2-pc-modules' installed
- Secure Boot is not supported in Linux.
- Boot disks that span multiple physical disks are not supported.
- Ensure that `/tmp` is mounted as read+write.
- Ensure that `/tmp` is mounted with the `exec` option. Verify that the `/tmp` directory is mounted in a way that allows you to run scripts and applications from it.

To verify that the `/tmp` directory is mounted without the `noexec` option, run this command: `sudo mount | grep '/tmp'`

If the result is similar to this example, it means that the issue exists in your OS:

```
$ sudo mount | grep '/tmp'
/dev/xvda1 on /tmp type ext4 (rw,noexec)
```

To fix and remove the *noexec* option from the mounted */tmp* directory, run this command: `sudo mount -o remount,exec /tmp`

This example illustrates the troubleshooting procedure:

```
ubuntu@Linux-1:~$ sudo mount | grep '/tmp'
/dev/xvda1 on /tmp type ext4 (rw,noexec)
ubuntu@Linux-1:~$ sudo mount -o remount,exec /tmp
ubuntu@Linux-1:~$ sudo mount | grep '/tmp'
/dev/xvda1 on /tmp type ext4 (rw)
```

- The AWS Transform MGN user needs to be either a root user or a user in the sudoers list.
- Ensure that the *dhclient* package is installed. If not, install the package using:

For Redhat/CentOS/Fedora/AmazonLinux:

```
sudo yum install dhclient
```

OR

```
sudo yum install dhcp-client
```

For Ubuntu/Debian:

```
sudo apt install isc-dhcp-client
```

For SUSE, check the [link](#) for the instructions to install "dhcp-client" package

- Verify that you have *kernel-devel/linux-headers* installed that are exactly the same version as the kernel you are running.

The version number of the kernel headers should be completely identical to the version number of the kernel. To handle this issue, follow these steps:

1. Identify the version of your running kernel.

```
uname -r
```

```
[root@ip-192-168-20-156 ~]# uname -r
4.14.177-107.254.amzn1.x86_64
[root@ip-192-168-20-156 ~]#
```

The `uname -r` output version should match the version of one of the installed kernel headers packages (`kernel-devel-<version number>` / `linux-headers-<version number>`).

2. Identify the version of your *kernel-devel/linux-headers*.

To identify the version of your installed `kernel-devel/linux-headers` packages, run this command:

On RHEL/CENTOS/Oracle/SUSE:

```
rpm -qa | grep kernel
```

```
[root@ip-192-168-20-156 ~]# rpm -qa |grep kernel
kernel-4.14.177-107.254.amzn1.x86_64
kernel-headers-4.14.181-108.257.amzn1.x86_64
kernel-devel-4.14.177-107.254.amzn1.x86_64
kernel-tools-4.14.181-108.257.amzn1.x86_64
[root@ip-192-168-20-156 ~]#
```

Note

This command looks for kernel related packages. The `kernel-devel` package is the specific package to look for.

On Debian/Ubuntu: `apt-cache search linux-headers`

```
ubuntu@Linux-1:~$ apt-cache search linux-headers
linux-headers-3.13.0-24 - Header files related to Linux kernel version
3.13.0
linux-headers-3.13.0-24-generic - Linux kernel headers for version 3.1
3.0 on 64 bit x86 SMP
linux-headers-3.13.0-24-lowlatency - Linux kernel headers for version
3.13.0 on 64 bit x86 SMP
```

3. Verify that the folder that contains the *kernel-devel/linux-headers* is not a symbolic link.

Sometimes, the content of the *kernel-devel/linux-headers*, which match the version of the kernel, is actually a symbolic link. In this case, you need to remove the link before installing the required package.


To verify that the folder that contains the *kernel-devel/linux-headers* is not a symbolic link, run this command:

On RHEL/CENTOS/Oracle:

```
ls -l /usr/src/kernels
```

On Debian/Ubuntu/SUSE:

```
ls -l /usr/src
```



```
ubuntu@Linux-1:~$ ls -l /usr/src
total 8
lrwxrwxrwx 1 root root 41 May 29 15:40 3.13.0-116-generic -> /usr/src/linux-headers-3.13.0-116-generic
drwxr-xr-x 24 root root 4096 Apr  5 20:43 linux-headers-3.13.0-116
drwxr-xr-x  7 root root 4096 Apr  5 20:43 linux-headers-3.13.0-116-generic
ubuntu@Linux-1:~$
```

In the above example, the results show that the actual **linux-headers-*** folders are not symbolic links.

4. [If a symbolic link exists] Delete the symbolic link.

If you found that the content of the *kernel-devel/linux-headers*, which match the version of the kernel, is a symbolic link, you need to delete the link. Run this command: `rm /usr/src/<LINK NAME>`

For example: `rm /usr/src/linux-headers-4.4.1`

5. Install the correct *kernel-devel/linux-headers* from the repositories.

If none of the already installed *kernel-devel/linux-headers* packages match your running kernel version, you need to install the matching package.

Note

You can have several kernel headers versions simultaneously on your OS, and you can therefore safely install new kernel headers packages in addition to your existing ones (without uninstalling the other versions of the package.) A new kernel headers

package does not impact the kernel, and does not overwrite older versions of the kernel headers.

 Note

For everything to work, you need to install a kernel headers package with the exact same version number of the running kernel.

To install the correct *kernel-devel/linux-headers*, run this command:

On RHEL/CENTOS/Oracle:

```
sudo yum install kernel-devel-`uname -r`
```

On Oracle with Unbreakable Enterprise Kernel:

```
sudo yum install kernel-uek-devel-`uname -r`
```

On Debian/Ubuntu:

```
sudo apt-get install linux-headers-`uname -r`
```

On SUSE:

```
sudo zypper install kernel-default-devel-`uname -r`
```

6. [If no matching package was found] Download the matching *kernel-devel/linux-headers* package.

If no matching package was found on the repositories configured on your server, you can download it manually from the Internet and then install it.

To download the matching *kernel-devel/linux-headers* package, navigate to these sites:

- [RHEL and Centos](#)
- [Oracle](#)
- [SUSE](#)
- [Debian](#)
- [Ubuntu](#)

If the *kernel-devel/linux-headers* packages are not available for the current running kernel version, consider upgrading the kernel to a version that has corresponding *kernel-devel* packages available. System administrators should validate that the appropriate *kernel-devel* packages are available before upgrading the kernel.

Windows installation requirements

Ensure that your source server operating system is supported. See [Supported Windows operating systems](#).

Ensure that your source server meets the agent installation hardware requirements, including:

- At least 2 GB of free disk space on the disk containing the "Program Files(x86)" directory
- Install all available Windows updates on the server.
- A graceful reboot from the OS menu or Windows CLI of a Windows source server does not trigger a rescan in MGN once the source server is restarted. Hard reboots, disk changes, and crashes trigger a rescan.
- Mount points must be assigned a drive letter to be recognized by AWS Transform MGN. A folder path is not recognized.

Operating systems supported by MGN

AWS Transform MGN supports replication of physical, virtual or cloud-based source servers for multiple versions of Windows and Linux operating systems.

Supported Windows operating systems

AWS Transform MGN allows replication of physical, virtual or cloud-based source servers to the AWS Cloud for multiple versions of Windows.

Note

Support deprecation notes

- **Windows 2003:** Effective February 15, 2026, this operating system is no longer supported.
- **Windows 2008:** Effective December 30, 2026, this operating system is no longer supported.
- **Windows 7:** Effective December 30, 2026, this operating system is no longer supported.

General Notes

- [Review the AWS Replication Agent installation requirements.](#)
- Windows source servers require a minimum of 2 GB of free disk space to launch a test or cutover instance.
- The WMI service must be activated to install the AWS Replication Agent.

These Windows operating systems are supported:

Operating system	Supported versions	Prerequisites and Limitations
Microsoft Windows Server 2025 64-bit		Requires .Net Framework version 4.5 or above to be installed by the end user.
Microsoft Windows Server 2022 64-bit		Requires .Net Framework version 4.5 or above to be installed by the end user.
Microsoft Windows Server 2019 64-bit		Requires .Net Framework version 4.5 or above to be installed by the end user.
Microsoft Windows Server 2016 64-bit		Requires .Net Framework version 4.5 or above to be installed by the end user.
Microsoft Windows 11 64-bit		Ensure that the auto sleep function is disabled. Data replication may be interrupted if the feature is activated.
Microsoft Windows 10 64-bit		Ensure that the auto sleep function in Windows 10 is disabled. Data replication may

Operating system	Supported versions	Prerequisites and Limitations
		be interrupted if the feature is activated.
Microsoft Windows Server 2012 This version has reached end of life. We recommend that you update to a more recent version.	64-bit and R2 64-bit	Requires .Net Framework version 4.5 or above to be installed by the end user.

Operating system	Supported versions	Prerequisites and Limitations
Microsoft Windows Server 2008 This version has reached end of life. We recommend that you update to a more recent version.	64-bit and R2 64-bit	- Windows Server 2008 requires .Net Framework version 3.5 to be installed by the end user. - Windows Server 2008 R2 requires .Net Framework version 4.5 or above to be installed by the end user. - Windows 2008 x64 requires SP2 and other Microsoft updates to support the SHA-2 signature of the AWS Replication Agent driver. - Windows Server 2008 R2 requires SP1 and the necessary KB updates for SHA-2 Code Signing Support to be installed prior to deploying the AWS Replication Agent. See 2019 SHA-2 Code Signing Support requirement for Windows and WSUS. - The AWS Replication Agent and agent installer requires a separate installer file, <code>AwsReplicationWindowsLegacyInstaller.exe</code> for end-of-life versions of Windows because they use older

Operating system	Supported versions	Prerequisites and Limitations
		versions of software components that cannot be upgraded. • Windows 2008 with GPT partitioned system drives are not supported. • Nitro instances can only be used with Windows Server 2008 R2 and upwards. Earlier versions are not supported. • The WMI service must be activated to install the AWS Replication Agent. • A shutdown (from the OS menu or Windows CLI) of a Windows source server triggers a rescan in AWS MGN once the source server is restarted.

Operating system	Supported versions	Prerequisites and Limitations
Microsoft Windows Server 2003 64-bit This version has reached end of life. We recommend that you update to a more recent version.		• Requires .Net Framework version 3.5 to be installed by the end user. • Does not support TLS 1.2, so you cannot download the AWS Replication Agent installer directly using the default browser. You must copy the file to the server using another method. • The AWS Replication Agent and agent installer requires a separate installer file, <code>AwsReplicationWindowsLegacyInstaller.exe</code> for end-of-life versions of Windows because they use older versions of software components that cannot be upgraded. • The WMI service must be activated to install the AWS Replication Agent. • Nitro instances can only be used with Windows Server 2008 R2 and upwards. Earlier versions are not supported. • A shutdown (from the OS menu or Windows CLI) of a Windows source server

Operating system	Supported versions	Prerequisites and Limitations
Microsoft Windows 7 64-bit This version has reached end of life. We recommend that you update to a more recent version.		triggers a rescan in AWS MGN once the source server is restarted. • The AWS Replication Agent and agent installer requires a separate installer file, <code>AwsReplicationWindowsLegacyInstaller.exe</code> for end-of-life versions of Windows because they use older versions of software components that cannot be upgraded. • Nitro instances can only be used with Windows Server 2008 R2 and upwards. Earlier versions are not supported. • A shutdown (from the OS menu or Windows CLI) of a Windows source server triggers a rescan in AWS MGN once the source server is restarted.

Supported Linux operating systems

Note

Support deprecation notes

- **Red Hat Enterprise Linux (RHEL) version 5.x and CentOS version 5.x:** Effective December 30, 2025, these operating systems are no longer supported.
- **Debian 6.x- 9.x:** Effective April 30, 2026, these operating systems are no longer supported.
- **Ubuntu 12.04:** Effective August 20, 2026, this operating system is no longer supported.
- **Oracle versions 6.x:** Effective August 28, 2026, these operating systems are no longer supported.
- **CentOS versions 6.x:** Effective August 28, 2026, these operating systems are no longer supported.
- **SLES versions 11.x:** Effective August 28, 2026, these operating systems are no longer supported.
- **CentOS 7-7.9:** Effective November 20, 2026, these operating systems are no longer supported.
- **Amazon Linux 1 (AL1):** Effective November 20, 2026, this operating system is no longer supported.
- **Ubuntu 14.04:** Effective December 20, 2026, this operating system is no longer supported.
- **Debian 10:** Effective December 30, 2026, this operating system is no longer supported.
- **Red Hat Enterprise Linux (RHEL) versions 6.x:** Effective December 30, 2026, these operating systems are no longer supported.
- **CentOS versions 8.x:** Effective December 30, 2026, these operating systems are no longer supported.

General Notes

- [Review the AWS Replication Agent installation requirements.](#)
- MGN does not support 32 bit versions of Linux.
- For source machines configured with LVM, on RHEL/Oracle version less than or equal to 9.4, make sure to update the lvm package to `lvm2-2.03.23-1.el9` or latest.
- Kernel version 4.9.256 is not supported. Agent installation fails on servers that run this kernel version.

- Kernel versions earlier than 2.6.18-164 are not supported by AWS Transform MGN. Therefore, servers that run these kernel versions cannot be replicated by AWS Transform MGN.

These Linux operating systems are supported:

Operating system	Supported versions	Prerequisites and Limitations
Amazon Linux	1, 2, 2023	• Amazon Linux 1 is only supported for AWS to AWS recovery. • Amazon Linux 1 is not supported in Canada West (Calgary). • Only agent-based replication is supported.
RHEL	6.0 to 9.8, 10, 10.1, 10.2	• For RHEL 8.x, a prerequisite is to run <code>\$ sudo yum install elfutils-libelf-devel</code> • Kernel versions 2.6.32-71 are not supported in RHEL 6.0 • The post-launch actions feature is not supported on RHEL 5.x and RHEL 6.x • Nitro instance types work with RHEL 7.4+ • AWS requires that servers running Red Hat Enterprise Linux (RHEL) must have Cloud Access (BYOL) licenses in order to be recovered to AWS. Note

Operating system	Supported versions	Prerequisites and Limitations
		that servers running RHEL Cloud Access Gold Images allow you to access AWS Red Hat Update Infrastructure (RHUI), Red Hat Satellite, or Red Hat Subscription Manager (RHSM). If you are using RHEL Cloud Access Gold Images, you are not able to access RHUI upon failover to AWS unless you link your AWS account to your Red Hat account via the Red Hat portal, and select the Gold image AMI in the launch template. • You must select an AWS provided RHEL AMI in the Launch Template for servers running Red Hat Enterprise Linux (RHEL) Pay as You Go (PAYG) images. This allows access to RHUI after migration. Note that usage of these images incurs Amazon EC2 charges for software and infrastructure per AWS Marketplace rates. • RHEL 8.3 and 8.4 are not supported in Mexico (Central) region.

Operating system	Supported versions	Prerequisites and Limitations
		• RHEL 8.3, 8.4, and 8.10 are not supported in Asia Pacific (New Zealand) and Asia Pacific (Taipei) regions. • FSx for ONTAP storage: If using FSx for ONTAP as the target storage type, pre-install iSCSI and multipath packages on the source server before migration: <code>sudo yum install -y iscsi-initiator-utils device-mapper-multipath</code> (or <code>dnf</code> for RHEL 8+). RHEL subscription credentials are tied to the source instance and are not valid on the migrated target.

Operating system	Supported versions	Prerequisites and Limitations
CentOS	7.0 to 8.0, Stream 9, Stream 10	- For Centos 8.x, a prerequisite is to run <code>\$ sudo yum install elfutils-libelf-devel</code> - Nitro instance types work with CentOS 7.4+ - CentOS Stream 9: these kernels were tested: 5.14.0-689.el9.x86_64 5.14.0-691.el9.x86_64 5.14.0-694.el9.x86_64 5.14.0-697.el9.x86_64 5.14.0-700.el9.x86_64 CentOS Stream is a rolling-release OS. Kernels outside this list are supported via local driver compilation. - CentOS Stream 10: these kernels are supported: 6.12.0-216.el10.x86_64 6.12.0-218.el10.x86_64 6.12.0-222.el10.x86_64 6.12.0-224.el10.x86_64 6.12.0-225.el10.x86_64 6.12.0-228.el10.x86_64 CentOS Stream is a rolling-release OS. Kernels outside this list are not supported.

Operating system	Supported versions	Prerequisites and Limitations
		• FSx for ONTAP storage: If using FSx for ONTAP as the target storage type, pre-install iSCSI and multipath packages on the source server before migration: <code>sudo yum install -y iscsi-initiator-utils device-mapper-multipath</code> .CentOS repository credentials are tied to the source instance and are not valid on the migrated target.

Operating system	Supported versions	Prerequisites and Limitations
Oracle Linux	7.0 to 7.8, 8.5 to 8.9, 9.0 to 9.4, 9.6, 9.7, and 10.1	- For Oracle Linux 8.x, a prerequisite is to run <code>\$ sudo yum install elfutils-libelf-devel</code> - Nitro instance types work with Oracle Linux 7.4+ - Oracle Linux 7.0 to 7.8 source servers must be running either Unbreakable Enterprise Kernel Release 3 or higher or a Red Hat Compatible Kernel. - Oracle Linux 8.5 to 8.9 (running either Unbreakable Enterprise Kernel Release 3 or higher or a Red Hat Compatible Kernel) – these UEK kernels were tested: 5.15.0-200.131.27.el9uek.x86_64 5.15.0-101.103.2.1.el9uek.x86_64 5.15.0-3.60.5.1.el9uek.x86_64 5.15.0-0.30.19.el9uek.x86_64 5.15.0-206.153.7.1.el8uek.x86_64 5.15.0-200.131.27.el8uek.x86_64

Operating system	Supported versions	Prerequisites and Limitations
		• 5.15.0-101.103.2.1.el8uek.x86_64 • 5.15.0-3.60.5.1.el8uek.x86_64 • 5.4.17-2136.314.6.3.el8uek.x86_64 • 5.4.17-2136.307.3.1.el8uek.x86_64 • 5.4.17-2136.300.7.el8uek.x86_64 • 4.18.0-372.32.1.0.1.el8_6.x86_64 • Oracle Linux 9.0 to 9.4 (running Unbreakable Enterprise Kernel Release 7 or Red Hat Compatible Kernel only) • Oracle Linux 9.6 and 9.7 (running Unbreakable Enterprise Kernel Release 8 or Red Hat Compatible Kernel only) • Oracle Linux 10.1 (running Unbreakable Enterprise Kernel Release 8 or Red Hat Compatible Kernel only) – these kernels are supported : • 6.12.0-100.28.2.el10uek.x86_64 - 6.12.0-105.51.5.el10uek.x86_64

Operating system	Supported versions	Prerequisites and Limitations
		6.12.0-124.8.1.el10_1 - 6.12.0-124.56.1.el10_1
Rocky Linux	8 to 9.8, 10, 10.1, 10.2	- For Rocky Linux 8.x, a prerequisite is to run <code>\$ sudo yum install elfutils-libelf-devel</code> - For UEFI-based Rocky Linux 10.2 systems without repository access, install <code>grub2-efi-x64-modules</code> before you run the agent installer.
SUSE Linux Enterprise Server	12 to 15 SP7	- The AWS Replication Agent is supported on SUSE Linux Enterprise Server (SLES) 12 and higher. FSx for ONTAP storage: If using FSx for ONTAP as the target storage type, pre-install iSCSI and multipath packages on the source server before migration: <code>sudo zypper install -y open-iscsi multipath-tools</code>. SLES subscription credentials are tied to the source instance and are not valid on the migrated target.

Operating system	Supported versions	Prerequisites and Limitations
Ubuntu	14.04 to 24.04	- Only Kernel 3.x or above are supported - Azure kernels are not supported as they are not compatible with the Amazon EC2 hardware. Ubuntu servers from Azure are required to switch the kernel to a standard kernel or the AWS tuned Ubuntu kernel 'linux-aws'.
Debian	10 to 11	Only Kernel 3.x or above are supported
AlmaLinux	8.6, 8.7, 8.8, 8.9, 8.10, 9.6, 9.7, 9.8, 10, 10.1, 10.2	Before you install the agent on AlmaLinux, complete the following prerequisites: - Run the following command with sudo privileges: <pre>\$ sudo yum install elfutils-libelf-devel</pre> - For UEFI-based AlmaLinux 9.8 and 10.2 systems without repository access, install <code>grub2-efi-x64-modules</code> before you run the agent installer.

Installing the AWS Replication Agent

You must install the AWS Replication Agent on each source server that you want to add to AWS Transform MGN. You need to reinstall the AWS Replication Agent after any major upgrade to the source server (see [Agent Related FAQ](#) for more info).

Topics

- [Generating the required AWS credentials](#)
- [Installing the AWS Replication Agent on Linux servers](#)
- [Installing the AWS Replication Agent on Windows servers](#)
- [Installing the Agent on a secured network](#)
- [Uninstalling the Agent](#)
- [Reinstalling the Agent](#)

Generating the required AWS credentials

In order to install the AWS Replication Agent, you must first generate the required AWS credentials.

Note

Temporary credentials have many advantages. You don't need to rotate them or revoke them when they're no longer needed, and they cannot be reused after they expire. You can specify for how long the credentials are valid, up to a maximum limit. Because they provide enhanced security, using temporary credentials is considered best practice and the recommended option.

Temporary credentials

The temporary credentials provided by AWS Transform MGN utilize a similar mechanism to the one used by [IAM Roles Anywhere](#).

To create temporary credentials, you need to:

1. [Create a new IAM Role](#) with the **AWSApplicationMigrationAgentInstallationPolicy** policy.
2. Request temporary security credentials [through AWS STS](#) through the [AssumeRole API](#).

An example of generating temporary credentials via AWS CLI can be found [here](#).

[Learn more about how temporary credentials work.](#)

Permanent credentials

Where possible, we recommend using temporary credentials instead of creating users who have long-term credentials such as passwords and access keys. However, there are specific use cases that require long-term credentials (for example, agentless snapshot based replications). In general, you should be able to use the same policy as recommended above in *Temporary Credentials*.

Installing the AWS Replication Agent on an Amazon EC2 instance

When installing an AWS Replication Agent on an Amazon EC2 instance (when the source server is in AWS Regions), you don't need to generate credentials. Instead, you can use an instance profile with the required IAM policy. This option can only be used when replicating a server within the same account.

- Go to the Amazon EC2 console and select your Amazon EC2 instance.
- From the top right-hand menu, select **Actions > Security > Modify IAM role**.
- Use a role that contains the [AWSApplicationMigrationServiceEc2InstancePolicy](#) policy.

If none exists, choose **Create new IAM role**, attach the policy and return to the Amazon EC2 console window.

- Select your new role from the drop-down list and choose **Update**.

Installing the AWS Replication Agent on Linux servers

Complete these steps to install the AWS Replication Agent on Linux source servers.

1. Ensure that the necessary service roles have been created by choosing the Reinitialize service permissions button on the AWS Transform MGN console's replication settings page. You must have the permissions necessary to create IAM roles in order for this operation to succeed.
2. Download the agent installer with the wget command on your Linux source server. This wget command downloads the Agent installer file - aws-replication-installer-init onto your server.

The Agent installer follows this format: `https://aws-application-migration-service-<region>.s3.<region>.amazonaws.com/latest/linux/aws-replication-installer-init`. Replace `<region>` with the AWS Region into which you are replicating.

This is an example of the full wget command for us-east-1:

```
wget -O ./aws-replication-installer-init https://aws-application-migration-service-us-east-1.s3.us-east-1.amazonaws.com/latest/linux/aws-replication-installer-init
```

The command line indicates when the installer has been successfully downloaded.

Important

- You need root privileges to run the Agent installer file on a Linux server. Alternatively, you can run the Agent Installer file with sudo permissions.
- If you need to validate the installer hash, the correct hash can be found here: `https://aws-application-migration-service-hashes-<region>.s3.<region>.amazonaws.com/latest/linux/aws-replication-installer-init.sha512` (replace `<region>` with the AWS Region into which you are replicating. For example, `us-east-1`:

`https://aws-application-migration-service-hashes-us-east-1.s3.us-east-1.amazonaws.com/latest/linux/aws-replication-installer-init.sha512`

- Replicating Amazon EC2 instances that were launched with marketplace product codes, is not supported.

Note

- The Linux installer creates the "aws-replication" group and "aws-replication" user within that group. The Agent runs within the context of the newly created user. Agent installation attempts to add the user to "sudoers". Installation fails if the Agent is unable to add the newly created "aws-replication" user to "sudoers".
- AWS Regions that are not opt-in also support the shorter installer path: `https://aws-application-migration-service-<region>.s3.amazonaws.com/latest/linux/aws-replication-installer-init`. Replace `<region>` with the AWS Region into which you are replicating.
- You can generate a custom installation command through the **Add servers** prompt. [Learn more about the Add servers prompt.](#)

This table contains the installer download link by supported AWS Region:

Region name	Region identity	Download Link
US East (Ohio)	us-east-2	IPv4 - https://aws-application-migration-service-us-east-2.s3.us-east-2.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-us-east-2.s3.dualstack.us-east-2.amazonaws.com/latest/linux/aws-replication-installer-init

Region name	Region identity	Download Link
US East (N. Virginia)	us-east-1	IPv4 - https://aws-application-migration-service-us-east-1.s3.us-east-1.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-us-east-1.s3.dualstack.us-east-1.amazonaws.com/latest/linux/aws-replication-installer-init
US West (N. California)	us-west-1	IPv4 - https://aws-application-migration-service-us-west-1.s3.us-west-1.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-us-west-1.s3.dualstack.us-west-1.amazonaws.com/latest/linux/aws-replication-installer-init

Region name	Region identity	Download Link
US West (Oregon)	us-west-2	IPv4 - https://aws-application-migration-service-us-west-2.s3.us-west-2.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-us-west-2.s3.dualstack.us-west-2.amazonaws.com/latest/linux/aws-replication-installer-init
Africa (Cape Town)	af-south-1	IPv4 - https://aws-application-migration-service-af-south-1.s3.af-south-1.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-af-south-1.s3.dualstack.af-south-1.amazonaws.com/latest/linux/aws-replication-installer-init

Region name	Region identity	Download Link
Asia Pacific (Thailand)	ap-southeast-7	IPv4 - https://aws-application-migration-service-ap-southeast-7.s3.ap-southeast-7.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-ap-southeast-7.s3.dualstack.ap-southeast-7.amazonaws.com/latest/linux/aws-replication-installer-init
Asia Pacific (New Zealand)	ap-southeast-6	IPv4 - https://aws-application-migration-service-ap-southeast-6.s3.ap-southeast-6.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-ap-southeast-6.s3.dualstack.ap-southeast-6.amazonaws.com/latest/linux/aws-replication-installer-init

Region name	Region identity	Download Link
Canada West (Calgary)	ca-west-1	IPv4 - https://aws-application-migration-service-ca-west-1.s3.ca-west-1.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-ca-west-1.s3.dualstack.ca-west-1.amazonaws.com/latest/linux/aws-replication-installer-init
Asia Pacific (Taipei)	ap-east-2	IPv4 - https://aws-application-migration-service-ap-east-2.s3.ap-east-2.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-ap-east-2.s3.dualstack.ap-east-2.amazonaws.com/latest/linux/aws-replication-installer-init

Region name	Region identity	Download Link
Mexico (Central)	mx-central-1	IPv4 - https://aws-application-migration-service-mx-central-1.s3.mx-central-1.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-mx-central-1.s3.dualstack.mx-central-1.amazonaws.com/latest/linux/aws-replication-installer-init
Asia Pacific (Hong Kong)	ap-east-1	IPv4 - https://aws-application-migration-service-ap-east-1.s3.ap-east-1.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-ap-east-1.s3.dualstack.ap-east-1.amazonaws.com/latest/linux/aws-replication-installer-init

Region name	Region identity	Download Link
Asia Pacific (Jakarta)	ap-southeast-3	IPv4 - https://aws-application-migration-service-ap-southeast-3.s3.ap-southeast-3.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-ap-southeast-3.s3.dualstack.ap-southeast-3.amazonaws.com/latest/linux/aws-replication-installer-init
Asia Pacific (Malaysia)	ap-southeast-5	IPv4 - https://aws-application-migration-service-ap-southeast-5.s3.ap-southeast-5.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-ap-southeast-5.s3.dualstack.ap-southeast-5.amazonaws.com/latest/linux/aws-replication-installer-init

Region name	Region identity	Download Link
Asia Pacific (Mumbai)	ap-south-1	IPv4 - https://aws-application-migration-service-ap-south-1.s3.ap-south-1.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-ap-south-1.s3.dualstack.ap-south-1.amazonaws.com/latest/linux/aws-replication-installer-init
Asia Pacific (Osaka)	ap-northeast-3	IPv4 - https://aws-application-migration-service-ap-northeast-3.s3.ap-northeast-3.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-ap-northeast-3.s3.dualstack.ap-northeast-3.amazonaws.com/latest/linux/aws-replication-installer-init

Region name	Region identity	Download Link
Asia Pacific (Seoul)	ap-northeast-2	IPv4 - https://aws-application-migration-service-ap-northeast-2.s3.ap-northeast-2.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-ap-northeast-2.s3.dualstack.ap-northeast-2.amazonaws.com/latest/linux/aws-replication-installer-init
Asia Pacific (Singapore)	ap-southeast-1	IPv4 - https://aws-application-migration-service-ap-southeast-1.s3.ap-southeast-1.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-ap-southeast-1.s3.dualstack.ap-southeast-1.amazonaws.com/latest/linux/aws-replication-installer-init

Region name	Region identity	Download Link
Asia Pacific (Sydney)	ap-southeast-2	IPv4 - https://aws-application-migration-service-ap-southeast-2.s3.ap-southeast-2.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-ap-southeast-2.s3.dualstack.ap-southeast-2.amazonaws.com/latest/linux/aws-replication-installer-init
Asia Pacific (Tokyo)	ap-northeast-1	IPv4 - https://aws-application-migration-service-ap-northeast-1.s3.ap-northeast-1.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-ap-northeast-1.s3.dualstack.ap-northeast-1.amazonaws.com/latest/linux/aws-replication-installer-init

Region name	Region identity	Download Link
Canada (Central)	ca-central-1	IPv4 - https://aws-application-migration-service-ca-central-1.s3.ca-central-1.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-ca-central-1.s3.dualstack.ca-central-1.amazonaws.com/latest/linux/aws-replication-installer-init
Europe (Frankfurt)	eu-central-1	IPv4 - https://aws-application-migration-service-eu-central-1.s3.eu-central-1.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-eu-central-1.s3.dualstack.eu-central-1.amazonaws.com/latest/linux/aws-replication-installer-init

Region name	Region identity	Download Link
Europe (Ireland)	eu-west-1	IPv4 - https://aws-application-migration-service-eu-west-1.s3.eu-west-1.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-eu-west-1.s3.dualstack.eu-west-1.amazonaws.com/latest/linux/aws-replication-installer-init
Europe (London)	eu-west-2	IPv4 - https://aws-application-migration-service-eu-west-2.s3.eu-west-2.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-eu-west-2.s3.dualstack.eu-west-2.amazonaws.com/latest/linux/aws-replication-installer-init

Region name	Region identity	Download Link
Europe (Milan)	eu-south-1	IPv4 - https://aws-application-migration-service-eu-south-1.s3.eu-south-1.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-eu-south-1.s3.dualstack.eu-south-1.amazonaws.com/latest/linux/aws-replication-installer-init
Europe (Paris)	eu-west-3	IPv4 - https://aws-application-migration-service-eu-west-3.s3.eu-west-3.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-eu-west-3.s3.dualstack.eu-west-3.amazonaws.com/latest/linux/aws-replication-installer-init

Region name	Region identity	Download Link
Europe (Stockholm)	eu-north-1	IPv4 - https://aws-application-migration-service-eu-north-1.s3.eu-north-1.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-eu-north-1.s3.dualstack.eu-north-1.amazonaws.com/latest/linux/aws-replication-installer-init
Middle East (Bahrain)	me-south-1	IPv4 - https://aws-application-migration-service-me-south-1.s3.me-south-1.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-me-south-1.s3.dualstack.me-south-1.amazonaws.com/latest/linux/aws-replication-installer-init

Region name	Region identity	Download Link
South America (São Paulo)	sa-east-1	IPv4 - https://aws-application-migration-service-sa-east-1.s3.sa-east-1.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-sa-east-1.s3.dualstack.sa-east-1.amazonaws.com/latest/linux/aws-replication-installer-init
Middle East (UAE)	me-central-1	IPv4 - https://aws-application-migration-service-me-central-1.s3.me-central-1.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-me-central-1.s3.dualstack.me-central-1.amazonaws.com/latest/linux/aws-replication-installer-init

Region name	Region identity	Download Link
Asia Pacific (Melbourne)	ap-southeast-4	IPv4 - https://aws-application-migration-service-ap-southeast-4.s3.ap-southeast-4.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-ap-southeast-4.s3.dualstack.ap-southeast-4.amazonaws.com/latest/linux/aws-replication-installer-init
Asia Pacific (Hyderabad)	ap-south-2	IPv4 - https://aws-application-migration-service-ap-south-2.s3.ap-south-2.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-ap-south-2.s3.dualstack.ap-south-2.amazonaws.com/latest/linux/aws-replication-installer-init

Region name	Region identity	Download Link
Europe (Zurich)	eu-central-2	IPv4 - https://aws-application-migration-service-eu-central-2.s3.eu-central-2.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-eu-central-2.s3.dualstack.eu-central-2.amazonaws.com/latest/linux/aws-replication-installer-init
Europe (Spain)	eu-south-2	IPv4 - https://aws-application-migration-service-eu-south-2.s3.eu-south-2.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-eu-south-2.s3.dualstack.eu-south-2.amazonaws.com/latest/linux/aws-replication-installer-init

Region name	Region identity	Download Link
Israel (Tel Aviv)	il-central-1	IPv4 - https://aws-application-migration-service-il-central-1.s3.il-central-1.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-il-central-1.s3.dualstack.il-central-1.amazonaws.com/latest/linux/aws-replication-installer-init
AWS GovCloud (US-East)	us-gov-east-1	IPv4 - https://aws-application-migration-service-us-gov-east-1.s3.us-gov-east-1.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-us-gov-east-1.s3.dualstack.us-gov-east-1.amazonaws.com/latest/linux/aws-replication-installer-init

Region name	Region identity	Download Link
AWS GovCloud (US-West)	us-gov-west-1	IPv4 - https://aws-application-migration-service-us-gov-west-1.s3.us-gov-west-1.amazonaws.com/latest/linux/aws-replication-installer-init Dual-stack - https://aws-application-migration-service-us-gov-west-1.s3.dualstack.us-gov-west-1.amazonaws.com/latest/linux/aws-replication-installer-init

3. [Generate the temporary credentials](#) that are required to install the AWS Replication Agent.

Important

When using [temporary credentials](#) (created using an IAM role instead of a user), you need to enter these parameters to the command prompt:

- AWS access key
- AWS secret access key
- AWS session token

The request to enter an AWS session token only appears if the system identifies that you are using temporary credentials. AWS access key for temporary credentials begins with the letters ASIA.

4. Once the agent installer has successfully downloaded, copy and input the installer command into the command line on your source server in order to run the installation script.

```
sudo chmod +x aws-replication-installer-init; sudo ./aws-replication-installer-init
```

You can choose to add the Region and required credential as parameters in the installation scripts:

- **--region** – The AWS Region in which the installer registers the source server.
- **--aws-access-key-id** – The AWS IAM Access Key used for authenticating the installing user. If this parameter is not provided, the installer prompts for it.
- **--aws-secret-access-key** – The AWS IAM Secret Access Key tied to the AWS IAM Access Key used for authenticating the installing user. If this parameter is not provided, the installer prompts for it.
- **--aws-session-token** – The session token is generated when using [temporary credentials](#) that are generated using AWS STS. If you use temporary credentials and do not provide this parameter, the installer prompts for it.

If you require additional customization, you can add a variety of parameters to the installation script in order to manipulate the way the agent is installed on your server. Add the parameters to the end of the installation script.

Available parameters include:

- **--no-prompt**

This parameter runs a silent installation.

- **--devices**

This parameter specifies which specific disks to replicate. The devices should be mentioned with comma separated, example `--devices="/dev/sda,/dev/sdb,/dev/sdc,/dev/sdd"`

- **--force-volumes**

This parameter must be used with the `--no-prompt` parameter. This parameter cancels the automatic detection of physical disks to replicate. You need to specify the exact disks to replicate using the `--devices` parameter (including the root disk, failure to specify the root disk causes replication to fail). This parameter should only be used as a troubleshooting tool if the `--devices` parameter fails to identify the disks correctly.

- **--tags**

Use this parameter to add resource tags to the source server. Use a space to separate each tag (for example: `--tags KEY=VALUE [KEY=VALUE ...]`)

Note

This flag may only be used when adding new source servers to MGN. You cannot use the `--tags` flag to modify tags of source servers that have already been added to MGN.

- `--s3-endpoint`

Use this parameter to specify a VPC endpoint you created for Amazon S3 if you do not wish to open your firewall ports to access the default Amazon S3 endpoint. [Learn more about installing the Agent on a blocked network.](#)

- `--user-provided-id`

This parameter allows you to provide a name to the source server that you are about to add, or identify a source server that needs to be updated. This identification is used by MGN to consistently recognize the server replication, and avoid duplication when [importing inventory](#) from a CSV file. Once provided for a server this parameter cannot be modified.

- `--endpoint`

Use this parameter:

- To specify the private link endpoint you created for AWS Transform MGN if you do not wish to open your firewall ports to access the default MGN endpoint. [Learn more about installing the Agent on a blocked network.](#)
- When using dual-stack, to specify a MGN dual-stack endpoint. You must also specify the `--dualstack` flag.

- `--no-replication`

By default after agent installation, the replication begins automatically. This attribute allows you to install the agent without immediately starting the replication. The 90-day free replication period excludes hours where the replication was stopped.

To start the replication post installation of replication agent using `--no-replication` attribute you can start replication by using the "Start Replication" option from Replication menu for the source server in the AWS MGN Dashboard or by using AWS CLI [start-replication](#)

- `--dualstack`

This parameter enables the agent to run in a dual-stack MGN configuration. When using this flag, you must also use the `--endpoint` flag to specify a MGN dual-stack endpoint.

- `--proxy-address`

Linux Installer only.

Use this parameter to configure the agent to use a specific proxy server: `--proxy-address https://PROXY:PORT/`. Make sure the proxy configuration has the trailing forward slash (/).

The installer confirms that the installation of the AWS Replication Agent has started.

```
root@ip-172-31-7-211:~# sudo chmod +x aws-replication-installer-init ./aws-replication-installer-init
The installation of the AWS Replication Agent has started.
```

5. The installer prompts you to enter your **AWS Region Name**, the **AWS Access Key ID**, the **AWS Secret Access Key**, and the **AWS Session Token** that you previously generated. Enter the complete AWS Region name (for example, eu-central-1) and the full credentials.

```
root@ip-172-31-7-211:~# sudo chmod +x aws-replication-installer-init ./aws-replication-installer-init
The installation of the AWS Replication Agent has started.
AWS Region Name: eu-west-1
AWS Access Key ID:
AWS Secret Access Key: *****
AWS Session Token: IGo7b1pZ1xK1YjE3R1////////wEaCXz1KVNc1QUM5JGMEQCI8g167WvpNsFt3dyPr:tDXanaTpjPCs0s1Dghd8g1BdzFFAdinw48oqc0j)+Knt15xb56uDF91o921wQgE1qnAgja////////8BBAAeADK3MTEAMj1QNTM5MyIMb1guXV5NurT4h2o8KVsBwNSM47a1N8TAgoN8yGJ7s1J0Sb++0wuh190ra+Qew3iF1zm8K68
```

If you want to install the Agent without answering the interactive prompts, you can pass your credentials to the installer through environment variables instead. We recommend that you use temporary credentials from AWS Security Token Service (AWS STS). First, set the `AWS_ACCESS_KEY_ID`, `AWS_SECRET_ACCESS_KEY`, and `AWS_SESSION_TOKEN` environment variables. Then, run the installer with the `sudo -E` command (to preserve the environment variables) and the `--no-prompt` option. For example:

```
$ export AWS_ACCESS_KEY_ID=AKIAIOSFODNN7EXAMPLE
$ export AWS_SECRET_ACCESS_KEY=wJa1rXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY
$ export AWS_SESSION_TOKEN=AQoDYXdzEJr////////wEa8AMDSomethingEXAMPLE
$ chmod +x aws-replication-installer-init
$ sudo -E ./aws-replication-installer-init --region us-east-1 --no-prompt
```

Note

- You can also pass the AWS Access Key ID and AWS Secret Access Key as command-line parameters. In the following example, replace `<region>` with the AWS Region into which you are replicating:

```
$ sudo chmod +x aws-replication-installer-init
```

```
$ sudo ./aws-replication-installer-init --region <region> --aws-access-
key-id AKIAIOSFODNN7EXAMPLE --aws-secret-access-key wJalrXUtnFEMI/K7MDENG/
bPxRfiCYEXAMPLEKEY
```

Warning

Credentials that you pass as command-line parameters are visible to other users on the source server through the process list (for example, through the `ps` command). To avoid exposing your credentials, use the environment variable method described in the previous step. If you do pass credentials as command-line parameters, use temporary credentials from AWS STS and rotate them after installation.

- If you do not enter these parameters as part of the installation script, you are prompted to enter them one by one as described in the previous step.
- The AWS Access Key ID and AWS Secret Access Key values are hidden when entered into the installer.

6. After you enter your credentials, the installer identifies volumes for replication. The installer displays the identified disks and prompts you to choose the disks you want to replicate.

```
root@ip-172-31-7-211:~# sudo chmod +x aws-replication-installer-init; ./aws-replication-installer-init
The installation of the AWS Replication Agent has started.
AWS Region Name: eu-west-1
AWS Access Key ID:
AWS Secret Access Key: *****
AWS Session Token: IqoN3lpZ21uXVjEJR//////////wAcXV1WVhc3QCMjJGMQCBj216WvpNaPtJdyR+xtDXanaTppCaf5oaNDqhcBjA1bzFYAd1nw4Boqo07j+KNCt5xb56u0P91o9Z1wQp81qAgja//////////BBEAEaDkMTEeKjQNTMSMyIMb1guXVSNUft4h2oBRv8w6SM4s7aNNtAgmNyD39sJUSb++0wuh19Dra+Qaw31P12mKX48
Identifying volumes for replication.
Choose the disks you want to replicate. Your disks are: /dev/xvdb,/dev/xvdc,/dev/xvda
To replicate some of the disks, type the path of the disks, separated with a comma (for example, /dev/sda,/dev/sdb). To replicate all disks, press Enter:
```

To replicate some of the disks, type the path of the disks, separated by a comma, as illustrated in the installer (such as: `/dev/sda`, `/dev/sdb`, and more). To replicate all of the disks, press **Enter**. The installer identifies the selected disks and prints their size.

```
Identifying volumes for replication.
Choose the disks you want to replicate. Your disks are: /dev/xvdb,/dev/xvdc,/dev/xvda
To replicate some of the disks, type the path of the disks, separated with a comma (for example, /dev/sda,/dev/sdb). To replicate all disks, press Enter:
Identified volume for replication: /dev/xvdb of size 10 GiB
Identified volume for replication: /dev/xvdc of size 1 GiB
Identified volume for replication: /dev/xvda of size 8 GiB
```

The installer confirms that all disks were successfully identified.

```
Identifying volumes for replication.
Choose the disks you want to replicate. Your disks are: /dev/xvdb,/dev/xvdc,/dev/xvda
To replicate some of the disks, type the path of the disks, separated with a comma (for example, /dev/sda,/dev/sdb). To replicate all disks, press Enter:
Identified volume for replication: /dev/xvdb of size 10 GiB
Identified volume for replication: /dev/xvdc of size 1 GiB
Identified volume for replication: /dev/xvda of size 8 GiB
```

Note

When identifying specific disks for replication, do not use apostrophes, brackets, or disk paths that do not exist. Type only existing disk paths. Each disk you selected for replication is displayed with the caption **Disk to replicate identified**. However, the displayed list of identified disks for replication may differ from the data you entered. This difference can be due to several reasons:

- The root disk of the source server is always replicated, whether you select it or not. Therefore, it always appears on the list of identified disks for replication.
- AWS Transform MGN replicates whole disks. Therefore, if you choose to replicate a partition, its entire disk appears on the list and is later replicated. If several partitions on the same disk are selected then that disk appears only once on the list.
- Incorrect disks may be chosen by accident. Ensure that the correct disks have been chosen.

Important

If disks are disconnected from a server, AWS Transform MGN can no longer replicate them, so they are removed from the list of replicated disks. When they are reconnected, the AWS Replication Agent cannot know that these were the same disks that were disconnected and therefore does not add them automatically. To add the disks after they are reconnected, rerun the AWS Replication Agent installer on the server.

Note that the returned disks need to be replicated from the beginning. Any disk size changes are automatically identified, but this also causes a resync. Perform a test after installing the Agent to ensure that the correct disks have been added.

7. After all of the disks that are to be replicated have been successfully identified, the installer downloads and installs the AWS Replication Agent on the source server.

```

root@ip-10-21-7-211:~# sudo chmod +x aws-replication-installer-init.sh ./aws-replication-installer-init.sh
The installation of the AWS Replication Agent has started.
AWS Region Name: eu-west-1
AWS Access Key ID:
AWS Secret Access Key: *****
AWS Session Token: IqoJb1pZ21uX2VjE3Rl//////////wEaCXv1WVhc3QMSjGMEQCIBg216NwvpaPtJdyE+z1DXanaTpJpCa5osNdphBgA1BdrFYAdin4BoqcoD7j+Kntt5xb56u09F91o921wqE1qgAqja//////////8BEAEaDDk3MTEwMjQ5MTM3My1Mb1guXV5Nuf4h2oBKv8BwN5M4e7ANRCAqmNyQ79eYUSB++0wuh19Dra+Qw31F12m8K64
Identifying volumes for replication.
Choose the disks you want to replicate. Your disks are: /dev/xvdb,/dev/xvdc,/dev/xvda
To replicate some of the disks, type the path of the disks, separated with a comma (for example, /dev/sda,/dev/sdb). To replicate all disks, press Enter:
Identified volume for replication: /dev/xvdb of size 10 GiB
Identified volume for replication: /dev/xvdc of size 1 GiB
Identified volume for replication: /dev/xvda of size 8 GiB
All volumes for replication were successfully identified.
Downloading the AWS Replication Agent onto the source server... Finished.
Installing the AWS Replication Agent onto the source server... Finished.

```

8. Once the AWS Replication Agent is installed, the server is added to the AWS Transform MGN console and undergoes the initial sync process. The installer provides you with the source server's ID.

```

root@ip-172-31-7-211:~# sudo chmod +x aws-replication-installer-init; aws-replication-installer-init
The installation of the AMS Replication Agent has started.
AWS Region Name: eu-west-1
AWS Access Key ID:
AWS Secret Access Key:
AWS Session Token: 10cbb3p21kx2vjy6H//////////wEACwZlWwncQpMcUgN8GQCjEpl6kNwvNpftJsydyvrtXtDdATpJCdsohNdgbdhSA1BdfFAdiaw4Boqpc07j7xNttbXsf6Hf9l0wCq1qAgJg//////////8BEAEaADok3MTEKwJQ2NTH5H1M1guxuV5NfUr4h2o8KvSaW0UM4e7AMk1AgmN8yD9j3yUSb+Dwhh190r4Qcw3P1I2mR56
Identifying volumes for replication.
Choose the disks you want to replicate. Your disks are /dev/xvdb, /dev/xvdc, /dev/xvdd
To replicate some of the disks, type the path of the disks, separated with a comma (for example, /dev/sda,/dev/sdb). To replicate all disks, press Enter:
Identified volume for replication: /dev/xvdb of size 10 GiB
Identified volume for replication: /dev/xvdc of size 1 GiB
Identified volume for replication: /dev/xvdd of size 6 GiB
You can also specify the disks by their size. For example, you can specify disks of size 10 GiB.
Download the AMS Replication Agent onto the source server... Finished.
Installing the AMS Replication Agent onto the source server... Finished.
Syncing the source server with the Application Migration Service Console... Finished.
Following is the source server ID: i-015f70p4wst
You now have 1 active source server out of a total quota of 200.
For more about increasing source server limit at https://docs.aws.amazon.com/mgn/latest/pg/MGN-service-limits.html
The AMS Replication Agent was successfully installed.

```

You can review this process in real time on the **Source servers** page. [Learn more about the initial sync process.](#)

Installing the AWS Replication Agent on Windows servers

Complete these steps to install the AWS Replication Agent on Windows source servers.

Ensure that the necessary service roles have been created by choosing the **Reinitialize service permissions** button on the AWS Transform MGN console replication settings page. You must have the permissions necessary to create IAM roles in order for this operation to succeed.

Download the agent installer (AWSReplicationWindowsInstaller.exe). Copy or distribute the downloaded agent installer to each Windows source server that you want to add to AWS Transform MGN.

The agent installer follows this format: `https://aws-application-migration-service-<region>.s3.<region>.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe`. Replace `<region>` with the AWS Region into which you are replicating.

This is an example of the installer link for us-east-1:

`https://aws-application-migration-service-us-east-1.s3.us-east-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe`

⚠ Important

- You need to run the agent installer file as an Administrator on each Windows server.
- If you need to validate the installer hash, the correct hash can be found here: <https://aws-application-migration-service-hashes->

`<region>.s3.<region>.amazonaws.com/latest/windows/
AwsReplicationWindowsInstaller.exe.sha512` (replace `<region>` with the AWS Region into which you are replicating, for example, `us-east-1`):

`https://aws-application-migration-service-hashes-us-east-1.s3.us-east-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512`

- We recommend using Windows PowerShell, which supports `ctrl+v` pasting, and not Windows Command Prompt (`cmd`), which does not.
- Replicating Amazon EC2 instances that were launched with marketplace product codes, is not supported.

Note

- AWS Regions that are not opt-in also support the shorter installer path: `https://aws-application-migration-service-<region>.s3.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe`. Replace `<region>` with the AWS Region into which you are replicating.
- You can generate a custom installation command through the **Add servers** prompt. [Learn more about the Add servers prompt.](#)
- Microsoft Windows Server versions 2003, 2003 R2, 2008, and 2008 R2 use a version of the AWS Replication Agent that is only valid for those versions - `AwsReplicationWindowsLegacyInstaller.exe`. DO NOT use this installer file to install the agent on any other OS types. You can generate an installer by following the steps outlined in the [Add servers actions prompt documentation](#) or directly download it from `https://aws-application-migration-service-<region>.s3.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe`. Replace `<region>` with the AWS Region into which you are replicating. If you need to validate the installer hash, the correct hash can be found here: `https://aws-application-migration-service-hashes-<region>.s3.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512` (replace `<region>` with the AWS Region into which you are replicating).
- Microsoft Windows Server 2012 uses a version of the AWS Replication Agent that is only valid for that version `AwsReplicationWindows2012LegacyInstaller.exe`.

DO NOT use this installer file to install the agent on any other OS types. You can download it from https://aws-application-migration-service-<REGION>.s3.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe. Replace <REGION> with the AWS Region into which you are replicating.

If you need to validate the installer hash, the correct hash can be found here: https://aws-application-migration-service-hashes-<region>.s3.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 (replace <region> with the AWS Region into which you are replicating).

AWS Replication Agent download URL for Windows for each supported AWS Region

Region name	Region identity	Download Link
US East (Ohio)	us-east-2	IPv4 - https://aws-application-migration-service-us-east-2.s3.us-east-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-us-east-2.s3.dualstack.us-east-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe
US East (N. Virginia)	us-east-1	IPv4 - https://aws-application-migration-service-us-east-1.s3.us-east-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe

Region name	Region identity	Download Link
		Dual-stack - https://aws-application-migration-service-us-east-1.s3.dualstack.us-east-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe
US West (N. California)	us-west-1	IPv4 - https://aws-application-migration-service-us-west-1.s3.us-west-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-us-west-1.s3.dualstack.us-west-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe
US West (Oregon)	us-west-2	IPv4 - https://aws-application-migration-service-us-west-2.s3.us-west-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-us-west-2.s3.dualstack.us-west-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe

Region name	Region identity	Download Link
Africa (Cape Town)	af-south-1	IPv4 - https://aws-application-migration-service-af-south-1.s3.af-south-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-af-south-1.s3.dualstack.af-south-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe
Asia Pacific (Hong Kong)	ap-east-1	IPv4 - https://aws-application-migration-service-ap-east-1.s3.ap-east-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-ap-east-1.s3.dualstack.ap-east-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe

Region name	Region identity	Download Link
Asia Pacific (Jakarta)	ap-southeast-3	IPv4 - https://aws-application-migration-service-ap-southeast-3.s3.ap-southeast-3.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-ap-southeast-3.s3.dualstack.ap-southeast-3.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe
Asia Pacific (Mumbai)	ap-south-1	IPv4 - https://aws-application-migration-service-ap-south-1.s3.ap-south-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-ap-south-1.s3.dualstack.ap-south-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe

Region name	Region identity	Download Link
Asia Pacific (Osaka)	ap-northeast-3	IPv4 - https://aws-application-migration-service-ap-northeast-3.s3.ap-northeast-3.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-ap-northeast-3.s3.dualstack.ap-northeast-3.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe
Asia Pacific (Seoul)	ap-northeast-2	IPv4 - https://aws-application-migration-service-ap-northeast-2.s3.ap-northeast-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-ap-northeast-2.s3.dualstack.ap-northeast-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe

Region name	Region identity	Download Link
Asia Pacific (Singapore)	ap-southeast-1	IPv4 - https://aws-application-migration-service-ap-southeast-1.s3.ap-southeast-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-ap-southeast-1.s3.dualstack.ap-southeast-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe
Asia Pacific (Sydney)	ap-southeast-2	IPv4 - https://aws-application-migration-service-ap-southeast-2.s3.ap-southeast-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-ap-southeast-2.s3.dualstack.ap-southeast-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe

Region name	Region identity	Download Link
Asia Pacific (Malaysia)	ap-southeast-5	IPv4 - https://aws-application-migration-service-ap-southeast-5.s3.ap-southeast-5.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-ap-southeast-5.s3.dualstack.ap-southeast-5.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe
Asia Pacific (Thailand)	ap-southeast-7	IPv4 - https://aws-application-migration-service-ap-southeast-7.s3.ap-southeast-7.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-ap-southeast-7.s3.dualstack.ap-southeast-7.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe

Region name	Region identity	Download Link
Asia Pacific (New Zealand)	ap-southeast-6	IPv4 - https://aws-application-migration-service-ap-southeast-6.s3.ap-southeast-6.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-ap-southeast-6.s3.dualstack.ap-southeast-6.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe
Canada West (Calgary)	ca-west-1	IPv4 - https://aws-application-migration-service-ca-west-1.s3.ca-west-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-ca-west-1.s3.dualstack.ca-west-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe

Region name	Region identity	Download Link
Asia Pacific (Taipei)	ap-east-2	IPv4 - https://aws-application-migration-service-ap-east-2.s3.ap-east-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-ap-east-2.s3.dualstack.ap-east-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe
Mexico (Central)	mx-central-1	IPv4 - https://aws-application-migration-service-mx-central-1.s3.mx-central-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-mx-central-1.s3.dualstack.mx-central-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe

Region name	Region identity	Download Link
Asia Pacific (Tokyo)	ap-northeast-1	IPv4 - https://aws-application-migration-service-ap-northeast-1.s3.ap-northeast-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-ap-northeast-1.s3.dualstack.ap-northeast-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe
Canada (Central)	ca-central-1	IPv4 - https://aws-application-migration-service-ca-central-1.s3.ca-central-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-ca-central-1.s3.dualstack.ca-central-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe

Region name	Region identity	Download Link
Europe (Frankfurt)	eu-central-1	IPv4 - https://aws-application-migration-service-eu-central-1.s3.eu-central-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-eu-central-1.s3.dualstack.eu-central-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe
Europe (Ireland)	eu-west-1	IPv4 - https://aws-application-migration-service-eu-west-1.s3.eu-west-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-eu-west-1.s3.dualstack.eu-west-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe

Region name	Region identity	Download Link
Europe (London)	eu-west-2	IPv4 - https://aws-application-migration-service-eu-west-2.s3.eu-west-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-eu-west-2.s3.dualstack.eu-west-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe
Europe (Milan)	eu-south-1	IPv4 - https://aws-application-migration-service-eu-south-1.s3.eu-south-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-eu-south-1.s3.dualstack.eu-south-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe

Region name	Region identity	Download Link
Europe (Paris)	eu-west-3	IPv4 - https://aws-application-migration-service-eu-west-3.s3.eu-west-3.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-eu-west-3.s3.dualstack.eu-west-3.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe
Europe (Stockholm)	eu-north-1	IPv4 - https://aws-application-migration-service-eu-north-1.s3.eu-north-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-eu-north-1.s3.dualstack.eu-north-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe

Region name	Region identity	Download Link
Middle East (Bahrain)	me-south-1	IPv4 - https://aws-application-migration-service-me-south-1.s3.me-south-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-me-south-1.s3.dualstack.me-south-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe
South America (São Paulo)	sa-east-1	IPv4 - https://aws-application-migration-service-sa-east-1.s3.sa-east-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-sa-east-1.s3.dualstack.sa-east-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe

Region name	Region identity	Download Link
Middle East (UAE)	me-central-1	IPv4 - https://aws-application-migration-service-me-central-1.s3.me-central-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-me-central-1.s3.dualstack.me-central-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe
Asia Pacific (Melbourne)	ap-southeast-4	IPv4 - https://aws-application-migration-service-ap-southeast-4.s3.ap-southeast-4.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-ap-southeast-4.s3.dualstack.ap-southeast-4.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe

Region name	Region identity	Download Link
Asia Pacific (Hyderabad)	ap-south-2	IPv4 - https://aws-application-migration-service-ap-south-2.s3.ap-south-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-ap-south-2.s3.dualstack.ap-south-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe
Europe (Zurich)	eu-central-2	IPv4 - https://aws-application-migration-service-eu-central-2.s3.eu-central-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-eu-central-2.s3.dualstack.eu-central-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe

Region name	Region identity	Download Link
Europe (Spain)	eu-south-2	IPv4 - https://aws-application-migration-service-eu-south-2.s3.eu-south-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-eu-south-2.s3.dualstack.eu-south-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe
Tel Aviv	il-central-1	IPv4 - https://aws-application-migration-service-il-central-1.s3.il-central-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-il-central-1.s3.dualstack.il-central-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe

Region name	Region identity	Download Link
AWS GovCloud (US-East)	us-gov-east-1	IPv4 - https://aws-application-migration-service-us-gov-east-1.s3.us-gov-east-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-us-gov-east-1.s3.dualstack.us-gov-east-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe
AWS GovCloud (US-West)	us-gov-west-1	IPv4 - https://aws-application-migration-service-us-gov-west-1.s3.us-gov-west-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe Dual-stack - https://aws-application-migration-service-us-gov-west-1.s3.dualstack.us-gov-west-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe

Validating the downloaded AWS Replication Agent installer for Windows.

Important

If you need to validate the installer hash, the correct hash is here:

```
https://aws-application-migration-service-hashes-
<REGION>.s3.<REGION>.amazonaws.com/latest/windows/
AwsReplicationWindowsInstaller.exe.sha512
```

Replace <REGION> with the AWS Region into which you are replicating, for example: us-east-1:

```
https://aws-application-migration-service-hashes-
us-east-1.s3.us-east-1.amazonaws.com/latest/windows/
AwsReplicationWindowsInstaller.exe.sha512
```

Region name	Region identity	SHA512 Hash Download Link
US East (N. Virginia)	us-east-1	IPv4 - https://aws-application-migration-service-hashes-us-east-1.s3.us-east-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-us-east-1.s3.dualstack.us-east-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512
US East (Ohio)	us-east-2	IPv4 - https://aws-application-migration-service-hashes-us-east-2.s3.us-east-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-us-east-2.s3.dualstack.us-east-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
		ualstack.us-east-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512
US West (N. California)	us-west-1	IPv4 - https://aws-application-migration-service-hashes-us-west-1.s3.us-west-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-us-west-1.s3.dualstack.us-west-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512
US West (Oregon)	us-west-2	IPv4 - https://aws-application-migration-service-hashes-us-west-2.s3.us-west-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-us-west-2.s3.dualstack.us-west-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Asia Pacific (Hong Kong)	ap-east-1	IPv4 - https://aws-application-migration-service-hashes-ap-east-1.s3.ap-east-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-east-1.s3.dualstack.ap-east-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512
Asia Pacific (Tokyo)	ap-northeast-1	IPv4 - https://aws-application-migration-service-hashes-ap-northeast-1.s3.ap-northeast-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-northeast-1.s3.dualstack.ap-northeast-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Asia Pacific (Seoul)	ap-northeast-2	IPv4 - https://aws-application-migration-service-hashes-ap-northeast-2.s3.ap-northeast-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-northeast-2.s3.dualstack.ap-northeast-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512
Asia Pacific (Osaka)	ap-northeast-3	IPv4 - https://aws-application-migration-service-hashes-ap-northeast-3.s3.ap-northeast-3.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-northeast-3.s3.dualstack.ap-northeast-3.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Asia Pacific (Singapore)	ap-southeast-1	IPv4 - https://aws-application-migration-service-hashes-ap-southeast-1.s3.ap-southeast-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-southeast-1.s3.dualstack.ap-southeast-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512
Asia Pacific (Sydney)	ap-southeast-2	IPv4 - https://aws-application-migration-service-hashes-ap-southeast-2.s3.ap-southeast-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-southeast-2.s3.dualstack.ap-southeast-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Asia Pacific (Jakarta)	ap-southeast-3	IPv4 - https://aws-application-migration-service-hashes-ap-southeast-3.s3.ap-southeast-3.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-southeast-3.s3.dualstack.ap-southeast-3.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512
Asia Pacific (Melbourne)	ap-southeast-4	IPv4 - https://aws-application-migration-service-hashes-ap-southeast-4.s3.ap-southeast-4.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-southeast-4.s3.dualstack.ap-southeast-4.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Asia Pacific (Malaysia)	ap-southeast-5	IPv4 - https://aws-application-migration-service-has-hes-ap-southeast-5.s3.ap-southeast-5.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-has-hes-ap-southeast-5.s3.dualstack.ap-southeast-5.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512
Asia Pacific (Thailand)	ap-southeast-7	IPv4 - https://aws-application-migration-service-has-hes-ap-southeast-7.s3.ap-southeast-7.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-has-hes-ap-southeast-7.s3.dualstack.ap-southeast-7.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Asia Pacific (New Zealand)	ap-southeast-6	IPv4 - https://aws-application-migration-service-hashes-ap-southeast-6.s3.ap-southeast-6.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-southeast-6.s3.dualstack.ap-southeast-6.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512
Canada West (Calgary)	ca-west-1	IPv4 - https://aws-application-migration-service-hashes-ca-west-1.s3.ca-west-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ca-west-1.s3.dualstack.ca-west-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Asia Pacific (Taipei)	ap-east-2	IPv4 - https://aws-application-migration-service-hashes-ap-east-2.s3.ap-east-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-east-2.s3.dualstack.ap-east-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512
Mexico (Central)	mx-central-1	IPv4 - https://aws-application-migration-service-hashes-mx-central-1.s3.mx-central-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-mx-central-1.s3.dualstack.mx-central-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Asia Pacific (Mumbai)	ap-south-1	IPv4 - https://aws-application-migration-service-has-hes-ap-south-1.s3.ap-south-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-has-hes-ap-south-1.s3.dualstack.ap-south-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512
Asia Pacific (Hyderabad)	ap-south-2	IPv4 - https://aws-application-migration-service-has-hes-ap-south-2.s3.ap-south-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-has-hes-ap-south-2.s3.dualstack.ap-south-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Europe (Frankfurt)	eu-central-1	IPv4 - https://aws-application-migration-service-hashes-eu-central-1.s3.eu-central-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-eu-central-1.s3.dualstack.eu-central-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512
Europe (Zurich)	eu-central-2	IPv4 - https://aws-application-migration-service-hashes-eu-central-2.s3.eu-central-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-eu-central-2.s3.dualstack.eu-central-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Europe (Stockholm)	eu-north-1	IPv4 - https://aws-application-migration-service-hashes-eu-north-1.s3.eu-north-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-eu-north-1.s3.dualstack.eu-north-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512
Europe (Milan)	eu-south-1	IPv4 - https://aws-application-migration-service-hashes-eu-south-1.s3.eu-south-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-eu-south-1.s3.dualstack.eu-south-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Europe (Spain)	eu-south-2	IPv4 - https://aws-application-migration-service-hashes-eu-south-2.s3.eu-south-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-eu-south-2.s3.dualstack.eu-south-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512
Europe (Ireland)	eu-west-1	IPv4 - https://aws-application-migration-service-hashes-eu-west-1.s3.eu-west-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-eu-west-1.s3.dualstack.eu-west-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Europe (London)	eu-west-2	IPv4 - https://aws-application-migration-service-has-hes-eu-west-2.s3.eu-west-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-has-hes-eu-west-2.s3.dualstack.eu-west-2.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512
Europe (Paris)	eu-west-3	IPv4 - https://aws-application-migration-service-has-hes-eu-west-3.s3.eu-west-3.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-has-hes-eu-west-3.s3.dualstack.eu-west-3.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Canada (Central)	ca-central-1	IPv4 - https://aws-application-migration-service-hashes-ca-central-1.s3.ca-central-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ca-central-1.s3.dualstack.ca-central-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512
Middle East (UAE)	me-central-1	IPv4 - https://aws-application-migration-service-hashes-me-central-1.s3.me-central-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-me-central-1.s3.dualstack.me-central-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Middle East (Bahrain)	me-south-1	IPv4 - https://aws-application-migration-service-hashes-me-south-1.s3.me-south-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-me-south-1.s3.dualstack.me-south-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512
South America (São Paulo)	sa-east-1	IPv4 - https://aws-application-migration-service-hashes-sa-east-1.s3.sa-east-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-sa-east-1.s3.dualstack.sa-east-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Africa (Cape Town)	af-south-1	IPv4 - https://aws-application-migration-service-has-hes-af-south-1.s3.af-south-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-has-hes-af-south-1.s3.dualstack.af-south-1.amazonaws.com/latest/windows/AwsReplicationWindowsInstaller.exe.sha512

AWS Replication Agent download URL for Windows versions 2003, 2003 R2, 2008, and 2008 R2 for each supported AWS Region

Region name	Region identity	Download Link
US East (N. Virginia)	us-east-1	IPv4 - https://aws-application-migration-service-us-east-1.s3.us-east-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-us-east-1.s3.dualstack.us-east-1.amazonaws.com/latest/windows_legacy/

Region name	Region identity	Download Link
		AwsReplicationWindowsLegacyInstaller.exe
US East (Ohio)	us-east-2	IPv4 - https://aws-application-migration-service-us-east-2.s3.us-east-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-us-east-2.s3.dualstack.us-east-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe
US West (N. California)	us-west-1	IPv4 - https://aws-application-migration-service-us-west-1.s3.us-west-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-us-west-1.s3.dualstack.us-west-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe

Region name	Region identity	Download Link
US West (Oregon)	us-west-2	IPv4 - https://aws-application-migration-service-us-west-2.s3.us-west-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-us-west-2.s3.dualstack.us-west-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe
Asia Pacific (Hong Kong)	ap-east-1	IPv4 - https://aws-application-migration-service-ap-east-1.s3.ap-east-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-east-1.s3.dualstack.ap-east-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe

Region name	Region identity	Download Link
Asia Pacific (Tokyo)	ap-northeast-1	IPv4 - https://aws-application-migration-service-ap-northeast-1.s3.ap-northeast-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-northeast-1.s3.dualstack.ap-northeast-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe
Asia Pacific (Seoul)	ap-northeast-2	IPv4 - https://aws-application-migration-service-ap-northeast-2.s3.ap-northeast-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-northeast-2.s3.dualstack.ap-northeast-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe

Region name	Region identity	Download Link
Asia Pacific (Osaka)	ap-northeast-3	IPv4 - https://aws-application-migration-service-ap-northeast-3.s3.ap-northeast-3.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-northeast-3.s3.dualstack.ap-northeast-3.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe
Asia Pacific (Singapore)	ap-southeast-1	IPv4 - https://aws-application-migration-service-ap-southeast-1.s3.ap-southeast-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-southeast-1.s3.dualstack.ap-southeast-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe

Region name	Region identity	Download Link
Asia Pacific (Sydney)	ap-southeast-2	IPv4 - https://aws-application-migration-service-ap-southeast-2.s3.ap-southeast-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-southeast-2.s3.dualstack.ap-southeast-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe
Asia Pacific (Jakarta)	ap-southeast-3	IPv4 - https://aws-application-migration-service-ap-southeast-3.s3.ap-southeast-3.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-southeast-3.s3.dualstack.ap-southeast-3.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe

Region name	Region identity	Download Link
Asia Pacific (Melbourne)	ap-southeast-4	IPv4 - https://aws-application-migration-service-ap-southeast-4.s3.ap-southeast-4.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-southeast-4.s3.dualstack.ap-southeast-4.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe
Asia Pacific (Malaysia)	ap-southeast-5	IPv4 - https://aws-application-migration-service-ap-southeast-5.s3.ap-southeast-5.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-southeast-5.s3.dualstack.ap-southeast-5.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe

Region name	Region identity	Download Link
Asia Pacific (Thailand)	ap-southeast-7	IPv4 - https://aws-application-migration-service-ap-southeast-7.s3.ap-southeast-7.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-southeast-7.s3.dualstack.ap-southeast-7.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe
Asia Pacific (New Zealand)	ap-southeast-6	IPv4 - https://aws-application-migration-service-ap-southeast-6.s3.ap-southeast-6.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-southeast-6.s3.dualstack.ap-southeast-6.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe

Region name	Region identity	Download Link
Canada West (Calgary)	ca-west-1	IPv4 - https://aws-application-migration-service-ca-west-1.s3.ca-west-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ca-west-1.s3.dualstack.ca-west-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe
Asia Pacific (Taipei)	ap-east-2	IPv4 - https://aws-application-migration-service-ap-east-2.s3.ap-east-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-east-2.s3.dualstack.ap-east-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe

Region name	Region identity	Download Link
Mexico (Central)	mx-central-1	IPv4 - https://aws-application-migration-service-mx-central-1.s3.mx-central-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-mx-central-1.s3.dualstack.mx-central-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe
Asia Pacific (Mumbai)	ap-south-1	IPv4 - https://aws-application-migration-service-ap-south-1.s3.ap-south-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-south-1.s3.dualstack.ap-south-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe

Region name	Region identity	Download Link
Asia Pacific (Hyderabad)	ap-south-2	IPv4 - https://aws-application-migration-service-ap-south-2.s3.ap-south-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-south-2.s3.dualstack.ap-south-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe
Europe (Frankfurt)	eu-central-1	IPv4 - https://aws-application-migration-service-eu-central-1.s3.eu-central-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-eu-central-1.s3.dualstack.eu-central-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe

Region name	Region identity	Download Link
Europe (Zurich)	eu-central-2	IPv4 - https://aws-application-migration-service-eu-central-2.s3.eu-central-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-eu-central-2.s3.dualstack.eu-central-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe
Europe (Stockholm)	eu-north-1	IPv4 - https://aws-application-migration-service-eu-north-1.s3.eu-north-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-eu-north-1.s3.dualstack.eu-north-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe

Region name	Region identity	Download Link
Europe (Milan)	eu-south-1	IPv4 - https://aws-application-migration-service-eu-south-1.s3.eu-south-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-eu-south-1.s3.dualstack.eu-south-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe
Europe (Spain)	eu-south-2	IPv4 - https://aws-application-migration-service-eu-south-2.s3.eu-south-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-eu-south-2.s3.dualstack.eu-south-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe

Region name	Region identity	Download Link
Europe (Ireland)	eu-west-1	IPv4 - https://aws-application-migration-service-eu-west-1.s3.eu-west-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-eu-west-1.s3.dualstack.eu-west-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe
Europe (London)	eu-west-2	IPv4 - https://aws-application-migration-service-eu-west-2.s3.eu-west-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-eu-west-2.s3.dualstack.eu-west-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe

Region name	Region identity	Download Link
Europe (Paris)	eu-west-3	IPv4 - https://aws-application-migration-service-eu-west-3.s3.eu-west-3.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-eu-west-3.s3.dualstack.eu-west-3.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe
Canada (Central)	ca-central-1	IPv4 - https://aws-application-migration-service-ca-central-1.s3.ca-central-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ca-central-1.s3.dualstack.ca-central-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe

Region name	Region identity	Download Link
Middle East (UAE)	me-central-1	IPv4 - https://aws-application-migration-service-me-central-1.s3.me-central-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-me-central-1.s3.dualstack.me-central-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe
Middle East (Bahrain)	me-south-1	IPv4 - https://aws-application-migration-service-me-south-1.s3.me-south-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-me-south-1.s3.dualstack.me-south-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe

Region name	Region identity	Download Link
South America (São Paulo)	sa-east-1	IPv4 - https://aws-application-migration-service-sa-east-1.s3.sa-east-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-sa-east-1.s3.dualstack.sa-east-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe
Africa (Cape Town)	af-south-1	IPv4 - https://aws-application-migration-service-af-south-1.s3.af-south-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe Dual-stack - https://aws-application-migration-service-af-south-1.s3.dualstack.af-south-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe

Validating the downloaded AWS Replication Agent installer for Windows versions 2003, 2003 R2, 2008, and 2008 R2.

Important

If you need to validate the installer hash, the correct hash is here:

`https://aws-application-migration-service-hashes-
<REGION>.s3.<REGION>.amazonaws.com/latest/windows_legacy/
AwsReplicationWindowsLegacyInstaller.exe.sha512`

Replace <REGION> with the AWS Region into which you are replicating, for example: us-east-1:

`https://aws-application-migration-service-hashes-us-
east-1.s3.us-east-1.amazonaws.com/latest/windows_legacy/
AwsReplicationWindowsLegacyInstaller.exe.sha512`

Region name	Region identity	SHA512 Hash Download Link
US East (N. Virginia)	us-east-1	IPv4 - <code>https://aws-application-migration-service-hashes-us-east-1.s3.us-east-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512</code> Dual-stack - <code>https://aws-application-migration-service-hashes-us-east-1.s3.dualstack.us-east-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512</code>
US East (Ohio)	us-east-2	IPv4 - <code>https://aws-application-migration-service-hashes-us-east-2.s3.us-east-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512</code>

Region name	Region identity	SHA512 Hash Download Link
		hes-us-east-2.s3.us-east-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-us-east-2.s3.dualstack.us-east-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512
US West (N. California)	us-west-1	IPv4 - https://aws-application-migration-service-hashes-us-west-1.s3.us-west-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-us-west-1.s3.dualstack.us-west-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
US West (Oregon)	us-west-2	IPv4 - https://aws-application-migration-service-hashes-us-west-2.s3.us-west-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-us-west-2.s3.dualstack.us-west-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512
Asia Pacific (Hong Kong)	ap-east-1	IPv4 - https://aws-application-migration-service-hashes-ap-east-1.s3.ap-east-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-east-1.s3.dualstack.ap-east-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Asia Pacific (Tokyo)	ap-northeast-1	IPv4 - https://aws-application-migration-service-hashes-ap-northeast-1.s3.ap-northeast-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-northeast-1.s3.dualstack.ap-northeast-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512
Asia Pacific (Seoul)	ap-northeast-2	IPv4 - https://aws-application-migration-service-hashes-ap-northeast-2.s3.ap-northeast-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-northeast-2.s3.dualstack.ap-northeast-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Asia Pacific (Osaka)	ap-northeast-3	IPv4 - https://aws-application-migration-service-hashes-ap-northeast-3.s3.ap-northeast-3.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-northeast-3.s3.dualstack.ap-northeast-3.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512
Asia Pacific (Singapore)	ap-southeast-1	IPv4 - https://aws-application-migration-service-hashes-ap-southeast-1.s3.ap-southeast-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-southeast-1.s3.dualstack.ap-southeast-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Asia Pacific (Sydney)	ap-southeast-2	IPv4 - https://aws-application-migration-service-hashes-ap-southeast-2.s3.ap-southeast-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-southeast-2.s3.dualstack.ap-southeast-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512
Asia Pacific (Jakarta)	ap-southeast-3	IPv4 - https://aws-application-migration-service-hashes-ap-southeast-3.s3.ap-southeast-3.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-southeast-3.s3.dualstack.ap-southeast-3.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Asia Pacific (Melbourne)	ap-southeast-4	IPv4 - https://aws-application-migration-service-hashes-ap-southeast-4.s3.ap-southeast-4.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-southeast-4.s3.dualstack.ap-southeast-4.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512
Asia Pacific (Malaysia)	ap-southeast-5	IPv4 - https://aws-application-migration-service-hashes-ap-southeast-5.s3.ap-southeast-5.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-southeast-5.s3.dualstack.ap-southeast-5.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Asia Pacific (Thailand)	ap-southeast-7	IPv4 - https://aws-application-migration-service-hashes-ap-southeast-7.s3.ap-southeast-7.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-southeast-7.s3.dualstack.ap-southeast-7.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512
Asia Pacific (New Zealand)	ap-southeast-6	IPv4 - https://aws-application-migration-service-hashes-ap-southeast-6.s3.ap-southeast-6.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-southeast-6.s3.dualstack.ap-southeast-6.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Canada West (Calgary)	ca-west-1	IPv4 - https://aws-application-migration-service-hashes-ca-west-1.s3.ca-west-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ca-west-1.s3.dualstack.ca-west-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512
Asia Pacific (Taipei)	ap-east-2	IPv4 - https://aws-application-migration-service-hashes-ap-east-2.s3.ap-east-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-east-2.s3.dualstack.ap-east-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Mexico (Central)	mx-central-1	IPv4 - https://aws-application-migration-service-hashes-mx-central-1.s3.mx-central-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-mx-central-1.s3.dualstack.mx-central-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512
Asia Pacific (Mumbai)	ap-south-1	IPv4 - https://aws-application-migration-service-hashes-ap-south-1.s3.ap-south-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-south-1.s3.dualstack.ap-south-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Asia Pacific (Hyderabad)	ap-south-2	IPv4 - https://aws-application-migration-service-hashes-ap-south-2.s3.ap-south-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-south-2.s3.dualstack.ap-south-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512
Europe (Frankfurt)	eu-central-1	IPv4 - https://aws-application-migration-service-hashes-eu-central-1.s3.eu-central-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-eu-central-1.s3.dualstack.eu-central-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Europe (Zurich)	eu-central-2	IPv4 - https://aws-application-migration-service-hashes-eu-central-2.s3.eu-central-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-eu-central-2.s3.dualstack.eu-central-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512
Europe (Stockholm)	eu-north-1	IPv4 - https://aws-application-migration-service-hashes-eu-north-1.s3.eu-north-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-eu-north-1.s3.dualstack.eu-north-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Europe (Milan)	eu-south-1	IPv4 - https://aws-application-migration-service-hashes-eu-south-1.s3.eu-south-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-eu-south-1.s3.dualstack.eu-south-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512
Europe (Spain)	eu-south-2	IPv4 - https://aws-application-migration-service-hashes-eu-south-2.s3.eu-south-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-eu-south-2.s3.dualstack.eu-south-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Europe (Ireland)	eu-west-1	IPv4 - https://aws-application-migration-service-hashes-eu-west-1.s3.eu-west-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-eu-west-1.s3.dualstack.eu-west-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512
Europe (London)	eu-west-2	IPv4 - https://aws-application-migration-service-hashes-eu-west-2.s3.eu-west-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-eu-west-2.s3.dualstack.eu-west-2.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Europe (Paris)	eu-west-3	IPv4 - https://aws-application-migration-service-hashes-eu-west-3.s3.eu-west-3.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-eu-west-3.s3.dualstack.eu-west-3.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512
Canada (Central)	ca-central-1	IPv4 - https://aws-application-migration-service-hashes-ca-central-1.s3.ca-central-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ca-central-1.s3.dualstack.ca-central-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Middle East (UAE)	me-central-1	IPv4 - https://aws-application-migration-service-hashes-me-central-1.s3.me-central-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-me-central-1.s3.dualstack.me-central-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512
Middle East (Bahrain)	me-south-1	IPv4 - https://aws-application-migration-service-hashes-me-south-1.s3.me-south-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-me-south-1.s3.dualstack.me-south-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
South America (São Paulo)	sa-east-1	IPv4 - https://aws-application-migration-service-hashes-sa-east-1.s3.sa-east-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-sa-east-1.s3.dualstack.sa-east-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512
Africa (Cape Town)	af-south-1	IPv4 - https://aws-application-migration-service-hashes-af-south-1.s3.af-south-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-af-south-1.s3.dualstack.af-south-1.amazonaws.com/latest/windows_legacy/AwsReplicationWindowsLegacyInstaller.exe.sha512

AWS Replication Agent download URL for Windows 2012 for each supported AWS Region

Region name	Region identity	Download Link
US East (N. Virginia)	us-east-1	IPv4 - https://aws-application-migration-service-us-east-1.s3.us-east-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-us-east-1.s3.dualstack.us-east-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe
US East (Ohio)	us-east-2	IPv4 - https://aws-application-migration-service-us-east-2.s3.us-east-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-us-east-2.s3.dualstack.us-east-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe

Region name	Region identity	Download Link
US West (N. California)	us-west-1	IPv4 - https://aws-application-migration-service-us-west-1.s3.us-west-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-us-west-1.s3.dualstack.us-west-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe
US West (Oregon)	us-west-2	IPv4 - https://aws-application-migration-service-us-west-2.s3.us-west-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-us-west-2.s3.dualstack.us-west-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe

Region name	Region identity	Download Link
Asia Pacific (Hong Kong)	ap-east-1	IPv4 - https://aws-application-migration-service-ap-east-1.s3.ap-east-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-east-1.s3.dualstack.ap-east-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe
Asia Pacific (Tokyo)	ap-northeast-1	IPv4 - https://aws-application-migration-service-ap-northeast-1.s3.ap-northeast-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-northeast-1.s3.dualstack.ap-northeast-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe

Region name	Region identity	Download Link
Asia Pacific (Seoul)	ap-northeast-2	IPv4 - https://aws-application-migration-service-ap-northeast-2.s3.ap-northeast-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-northeast-2.s3.dualstack.ap-northeast-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe

Region name	Region identity	Download Link
Asia Pacific (Osaka)	ap-northeast-3	IPv4 - https://aws-application-migration-service-ap-northeast-3.s3.ap-northeast-3.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-northeast-3.s3.dualstack.ap-northeast-3.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe

Region name	Region identity	Download Link
Asia Pacific (Singapore)	ap-southeast-1	IPv4 - https://aws-application-migration-service-ap-southeast-1.s3.ap-southeast-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-southeast-1.s3.dualstack.ap-southeast-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe

Region name	Region identity	Download Link
Asia Pacific (Sydney)	ap-southeast-2	IPv4 - https://aws-application-migration-service-ap-southeast-2.s3.ap-southeast-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-southeast-2.s3.dualstack.ap-southeast-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe

Region name	Region identity	Download Link
Asia Pacific (Jakarta)	ap-southeast-3	IPv4 - https://aws-application-migration-service-ap-southeast-3.s3.ap-southeast-3.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-southeast-3.s3.dualstack.ap-southeast-3.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe

Region name	Region identity	Download Link
Asia Pacific (Melbourne)	ap-southeast-4	IPv4 - https://aws-application-migration-service-ap-southeast-4.s3.ap-southeast-4.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-southeast-4.s3.dualstack.ap-southeast-4.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe

Region name	Region identity	Download Link
Asia Pacific (Malaysia)	ap-southeast-5	IPv4 - https://aws-application-migration-service-ap-southeast-5.s3.ap-southeast-5.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-southeast-5.s3.dualstack.ap-southeast-5.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe

Region name	Region identity	Download Link
Asia Pacific (Thailand)	ap-southeast-7	IPv4 - https://aws-application-migration-service-ap-southeast-7.s3.ap-southeast-7.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-southeast-7.s3.dualstack.ap-southeast-7.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe

Region name	Region identity	Download Link
Asia Pacific (New Zealand)	ap-southeast-6	IPv4 - https://aws-application-migration-service-ap-southeast-6.s3.ap-southeast-6.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-southeast-6.s3.dualstack.ap-southeast-6.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe
Canada West (Calgary)	ca-west-1	IPv4 - https://aws-application-migration-service-ca-west-1.s3.ca-west-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ca-west-1.s3.dualstack.ca-west-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe

Region name	Region identity	Download Link
Asia Pacific (Taipei)	ap-east-2	IPv4 - https://aws-application-migration-service-ap-east-2.s3.ap-east-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-east-2.s3.dualstack.ap-east-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe
Mexico (Central)	mx-central-1	IPv4 - https://aws-application-migration-service-mx-central-1.s3.mx-central-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-mx-central-1.s3.dualstack.mx-central-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe

Region name	Region identity	Download Link
Asia Pacific (Mumbai)	ap-south-1	IPv4 - https://aws-application-migration-service-ap-south-1.s3.ap-south-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-south-1.s3.dualstack.ap-south-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe

Region name	Region identity	Download Link
Asia Pacific (Hyderabad)	ap-south-2	IPv4 - https://aws-application-migration-service-ap-south-2.s3.ap-south-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ap-south-2.s3.dualstack.ap-south-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe
Europe (Frankfurt)	eu-central-1	IPv4 - https://aws-application-migration-service-eu-central-1.s3.eu-central-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-eu-central-1.s3.dualstack.eu-central-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe

Region name	Region identity	Download Link
Europe (Zurich)	eu-central-2	IPv4 - https://aws-application-migration-service-eu-central-2.s3.eu-central-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-eu-central-2.s3.dualstack.eu-central-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe
Europe (Stockholm)	eu-north-1	IPv4 - https://aws-application-migration-service-eu-north-1.s3.eu-north-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-eu-north-1.s3.dualstack.eu-north-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe

Region name	Region identity	Download Link
Europe (Milan)	eu-south-1	IPv4 - https://aws-application-migration-service-eu-south-1.s3.eu-south-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-eu-south-1.s3.dualstack.eu-south-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe

Region name	Region identity	Download Link
Europe (Spain)	eu-south-2	IPv4 - https://aws-application-migration-service-eu-south-2.s3.eu-south-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-eu-south-2.s3.dualstack.eu-south-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe
Europe (Ireland)	eu-west-1	IPv4 - https://aws-application-migration-service-eu-west-1.s3.eu-west-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-eu-west-1.s3.dualstack.eu-west-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe

Region name	Region identity	Download Link
Europe (London)	eu-west-2	IPv4 - https://aws-application-migration-service-eu-west-2.s3.eu-west-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-eu-west-2.s3.dualstack.eu-west-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe
Europe (Paris)	eu-west-3	IPv4 - https://aws-application-migration-service-eu-west-3.s3.eu-west-3.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-eu-west-3.s3.dualstack.eu-west-3.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe

Region name	Region identity	Download Link
Canada (Central)	ca-central-1	IPv4 - https://aws-application-migration-service-ca-central-1.s3.ca-central-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-ca-central-1.s3.dualstack.ca-central-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe
Middle East (UAE)	me-central-1	IPv4 - https://aws-application-migration-service-me-central-1.s3.me-central-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-me-central-1.s3.dualstack.me-central-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe

Region name	Region identity	Download Link
Middle East (Bahrain)	me-south-1	IPv4 - https://aws-application-migration-service-me-south-1.s3.me-south-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-me-south-1.s3.dualstack.me-south-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe
South America (São Paulo)	sa-east-1	IPv4 - https://aws-application-migration-service-sa-east-1.s3.sa-east-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-sa-east-1.s3.dualstack.sa-east-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe

Region name	Region identity	Download Link
Africa (Cape Town)	af-south-1	IPv4 - https://aws-application-migration-service-af-south-1.s3.af-south-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe Dual-stack - https://aws-application-migration-service-af-south-1.s3.dualstack.af-south-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe

Validating the downloaded AWS Replication Agent installer for Windows 2012.

Important

If you need to validate the installer hash, the correct hash is here:

```
https://aws-application-migration-service-
hashes-<REGION>.s3.<REGION>.amazonaws.com/
latest/windows_legacy/windows_2012_legacy/
AwsReplicationWindows2012LegacyInstaller.exe.sha512
```

Replace <REGION> with the AWS Region into which you are replicating, for example: us-east-1:

```
https://aws-application-migration-service-hashes-us-east-1.s3.us-
east-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/
AwsReplicationWindows2012LegacyInstaller.exe.sha512
```

Region name	Region identity	SHA512 Hash Download Link
US East (N. Virginia)	us-east-1	IPv4 - https://aws-application-migration-service-hashes-us-east-1.s3.us-east-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-us-east-1.s3.dualstack.us-east-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512
US East (Ohio)	us-east-2	IPv4 - https://aws-application-migration-service-hashes-us-east-2.s3.us-east-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-us-east-2.s3.dualstack.us-east-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
US West (N. California)	us-west-1	IPv4 - https://aws-application-migration-service-hashes-us-west-1.s3.us-west-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-us-west-1.s3.dualstack.us-west-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512
US West (Oregon)	us-west-2	IPv4 - https://aws-application-migration-service-hashes-us-west-2.s3.us-west-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-us-west-2.s3.dualstack.us-west-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Asia Pacific (Hong Kong)	ap-east-1	IPv4 - https://aws-application-migration-service-hashes-ap-east-1.s3.ap-east-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-east-1.s3.dualstack.ap-east-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Asia Pacific (Tokyo)	ap-northeast-1	IPv4 - https://aws-application-migration-service-hashes-ap-northeast-1.s3.ap-northeast-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-northeast-1.s3.dualstack.ap-northeast-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Asia Pacific (Seoul)	ap-northeast-2	IPv4 - https://aws-application-migration-service-hashes-ap-northeast-2.s3.ap-northeast-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-northeast-2.s3.dualstack.ap-northeast-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Asia Pacific (Osaka)	ap-northeast-3	IPv4 - https://aws-application-migration-service-hashes-ap-northeast-3.s3.ap-northeast-3.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-northeast-3.s3.dualstack.ap-northeast-3.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Asia Pacific (Singapore)	ap-southeast-1	IPv4 - https://aws-application-migration-service-has-hes-ap-southeast-1.s3.ap-southeast-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-has-hes-ap-southeast-1.s3.dualstack.ap-southeast-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Asia Pacific (Sydney)	ap-southeast-2	IPv4 - https://aws-application-migration-service-hashes-ap-southeast-2.s3.ap-southeast-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-southeast-2.s3.dualstack.ap-southeast-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Asia Pacific (Jakarta)	ap-southeast-3	IPv4 - https://aws-application-migration-service-hashes-ap-southeast-3.s3.ap-southeast-3.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-southeast-3.s3.dualstack.ap-southeast-3.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Asia Pacific (Melbourne)	ap-southeast-4	IPv4 - https://aws-application-migration-service-hashes-ap-southeast-4.s3.ap-southeast-4.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-southeast-4.s3.dualstack.ap-southeast-4.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Asia Pacific (Malaysia)	ap-southeast-5	IPv4 - https://aws-application-migration-service-hashes-ap-southeast-5.s3.ap-southeast-5.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-southeast-5.s3.dualstack.ap-southeast-5.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Asia Pacific (Thailand)	ap-southeast-7	IPv4 - https://aws-application-migration-service-hashes-ap-southeast-7.s3.ap-southeast-7.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-southeast-7.s3.dualstack.ap-southeast-7.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Asia Pacific (New Zealand)	ap-southeast-6	IPv4 - https://aws-application-migration-service-hashes-ap-southeast-6.s3.ap-southeast-6.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-southeast-6.s3.dualstack.ap-southeast-6.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Canada West (Calgary)	ca-west-1	IPv4 - https://aws-application-migration-service-hashes-ca-west-1.s3.ca-west-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ca-west-1.s3.dualstack.ca-west-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512
Asia Pacific (Taipei)	ap-east-2	IPv4 - https://aws-application-migration-service-hashes-ap-east-2.s3.ap-east-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-east-2.s3.dualstack.ap-east-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Mexico (Central)	mx-central-1	IPv4 - https://aws-application-migration-service-hashes-mx-central-1.s3.mx-central-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-mx-central-1.s3.dualstack.mx-central-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Asia Pacific (Mumbai)	ap-south-1	IPv4 - https://aws-application-migration-service-hashes-ap-south-1.s3.ap-south-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-south-1.s3.dualstack.ap-south-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Asia Pacific (Hyderabad)	ap-south-2	IPv4 - https://aws-application-migration-service-hashes-ap-south-2.s3.ap-south-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ap-south-2.s3.dualstack.ap-south-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Europe (Frankfurt)	eu-central-1	IPv4 - https://aws-application-migration-service-hashes-eu-central-1.s3.eu-central-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-eu-central-1.s3.dualstack.eu-central-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Europe (Zurich)	eu-central-2	IPv4 - https://aws-application-migration-service-hashes-eu-central-2.s3.eu-central-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-eu-central-2.s3.dualstack.eu-central-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Europe (Stockholm)	eu-north-1	IPv4 - https://aws-application-migration-service-hashes-eu-north-1.s3.eu-north-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-eu-north-1.s3.dualstack.eu-north-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Europe (Milan)	eu-south-1	IPv4 - https://aws-application-migration-service-hashes-eu-south-1.s3.eu-south-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-eu-south-1.s3.dualstack.eu-south-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Europe (Spain)	eu-south-2	IPv4 - https://aws-application-migration-service-hashes-eu-south-2.s3.eu-south-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-eu-south-2.s3.dualstack.eu-south-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Europe (Ireland)	eu-west-1	IPv4 - https://aws-application-migration-service-hashes-eu-west-1.s3.eu-west-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-eu-west-1.s3.dualstack.eu-west-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512
Europe (London)	eu-west-2	IPv4 - https://aws-application-migration-service-hashes-eu-west-2.s3.eu-west-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-eu-west-2.s3.dualstack.eu-west-2.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Europe (Paris)	eu-west-3	IPv4 - https://aws-application-migration-service-hashes-eu-west-3.s3.eu-west-3.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-eu-west-3.s3.dualstack.eu-west-3.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Canada (Central)	ca-central-1	IPv4 - https://aws-application-migration-service-hashes-ca-central-1.s3.ca-central-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-ca-central-1.s3.dualstack.ca-central-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Middle East (UAE)	me-central-1	IPv4 - https://aws-application-migration-service-hashes-me-central-1.s3.me-central-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-me-central-1.s3.dualstack.me-central-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Middle East (Bahrain)	me-south-1	IPv4 - https://aws-application-migration-service-hashes-me-south-1.s3.me-south-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-me-south-1.s3.dualstack.me-south-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
South America (São Paulo)	sa-east-1	IPv4 - https://aws-application-migration-service-hashes-sa-east-1.s3.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-sa-east-1.s3.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Region name	Region identity	SHA512 Hash Download Link
Africa (Cape Town)	af-south-1	IPv4 - https://aws-application-migration-service-hashes-af-south-1.s3.af-south-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512 Dual-stack - https://aws-application-migration-service-hashes-af-south-1.s3.dualstack.af-south-1.amazonaws.com/latest/windows_legacy/windows_2012_legacy/AwsReplicationWindows2012LegacyInstaller.exe.sha512

Installation steps

Complete these steps to install the AWS Replication Agent on Windows source servers.

1. [Generate the temporary credentials](#) that are required to install the AWS Replication Agent.

Important

When using [temporary credentials](#) (created using an IAM role instead of a user), you need to enter these parameters to the command prompt:

- AWS access key
- AWS secret access key
- AWS session token

The request to enter an AWS session token only appears if the system identifies that you are using temporary credentials. AWS access key for temporary credentials begins with the letters ASIA.

2. Run the agent installer file – `AWSReplicationWindowsInstaller.exe` – as an Administrator. The CMD opens.

To run the installer with the default settings, enter your **AWS Region Name**, the **AWS Access Key ID** and the **AWS Secret Access Key**, and the **AWS Session Token** as described in the next step.

If you require additional customization, you can add a variety of parameters to the installation script in order to manipulate the way the Agent is installed on your server. Add the parameters to the end of the installation script.

- `--no-prompt`

This parameter runs a silent installation.

- `--devices`

This parameter specifies which specific disks to replicate.

- `--force-volumes`

This parameter must be used with the `--no-prompt` parameter. This parameter cancels the automatic detection of physical disks to replicate. You need to specify the exact disks to replicate using the `--devices` parameter (including the root disk, failure to specify the root disk causes replication to fail). This parameter should only be used as a troubleshooting tool if the `--devices` parameter fails to identify the disks correctly.

- `--tags`

Use this parameter to add resource tags to the source server. Use a space to separate each tag (for example: `--tags KEY=VALUE [KEY=VALUE ...]`)

Note

This flag may only be used when adding new source servers to MGN. You cannot use the `--tags` flag to modify tags of source servers that have already been added to MGN.

- `--s3-endpoint`

Use this parameter to specify a VPC endpoint you created for Amazon S3 if you do not wish to open your firewall ports to access the default Amazon S3 endpoint. [Learn more about installing the Agent on a blocked network.](#)

- `--user-provided-id`

This parameter allows you to provide a name to the source server that you are about to add, or identify a source server that needs to be updated. This identification is used by MGN to consistently recognize the server replication, and avoid duplication when [importing inventory](#) from a CSV file. Once provided for a server this parameter cannot be modified.

- `--endpoint`

Use this parameter to specify the Private Link endpoint you created for AWS Transform MGN if you do not wish to open your firewall ports to access the default MGN endpoint. [Learn more about installing the agent on a blocked network.](#)


- `--no-replication`

By default after agent installation, the replication begins automatically. This attribute allows you to install the agent without immediately starting the replication. The 90-day free replication period excludes hours where the replication was stopped.


- `--dualstack`

This parameter enables the agent to run in a dual-stack MGN configuration.

The installer confirms that the installation of the AWS Replication Agent has started.



3. The installer prompts you to enter your **AWS Region Name**, the **AWS Access Key ID**, the **AWS Secret Access Key** (and the **AWS Session Token** if appropriate) that you previously generated. Enter the complete AWS Region name (for example: eu-central-1), and the full AWS Access Key ID and AWS Secret Access Key.



If you want to install the Agent without answering the interactive prompts, you can pass your credentials to the installer through environment variables instead. We recommend that you

use temporary credentials from AWS Security Token Service (AWS STS). Open an elevated (Administrator) PowerShell session, set the `AWS_ACCESS_KEY_ID`, `AWS_SECRET_ACCESS_KEY`, and `AWS_SESSION_TOKEN` environment variables, and then run the installer with the `--no-prompt` option. The following PowerShell example sets these variables and runs the installer:

```
PS C:\> $env:AWS_ACCESS_KEY_ID="AKIAIOSFODNN7EXAMPLE"
PS C:\> $env:AWS_SECRET_ACCESS_KEY="wJalrXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY"
PS C:\> $env:AWS_SESSION_TOKEN="AQoDYXdzEJr////////wEa8AMDSomethingEXAMPLE"
PS C:\> .\AwsReplicationWindowsInstaller.exe --region us-east-1 --no-prompt
```

Note

- You can also pass the AWS Access Key ID and AWS Secret Access Key as command-line parameters. In the following example, replace `regionname` with the AWS Region into which you are replicating:


```
C:\> AwsReplicationWindowsInstaller.exe --region regionname --aws-access-  
key-id AKIAIOSFODNN7EXAMPLE --aws-secret-access-key wJalrXUtnFEMI/K7MDENG/  
bPxrFiCYEXAMPLEKEY
```

Warning

Credentials that you pass as command-line parameters are visible to other users on the source server. They can view these credentials through the process list (for example, through the `Get-CimInstance Win32_Process` PowerShell command). To avoid exposing your credentials, use the environment variable method described in the previous step. If you do pass credentials as command-line parameters, use temporary credentials from AWS STS and rotate them after you install the Agent.


- If you do not enter these parameters as part of the installation script, you are prompted to enter them one by one as described in the previous step.
- The AWS Access Key ID and AWS Secret Access Key values are hidden when entered into the installer.

4. Once you have entered your credentials, the installer verifies that the source server has enough free disk space for Agent installation and identifies volumes for replication. The installer displays the identified disks and prompts you to choose the disks you want to replicate.




```
C:\Users\Administrator\Downloads\AwsReplicationWindows\Installer.exe
The installation of the AWS Replication Agent has started.
AWS Region Name: eu-west-1
AWS Access Key ID:
AWS Secret Access Key:
Verifying that the source server has enough free disk space to install the AWS Replication Agent.
(a minimum of 2 GB of free disk space is required)
Identifying volumes for replication.
Choose the disks you want to replicate. Your disks are: c:
```

To replicate some of the disks, type the path of the disks, separated by a comma, as illustrated in the installer (for example: C: or D:). To replicate all of the disks, press **Enter**. The installer identifies the selected disks and prints their size.



```
C:\Users\Administrator\Downloads\AwsReplicationWindows\Installer.exe
The installation of the AWS Replication Agent has started.
AWS Region Name: eu-west-1
AWS Access Key ID:
AWS Secret Access Key:
Verifying that the source server has enough free disk space to install the AWS Replication Agent.
(a minimum of 2 GB of free disk space is required)
Identifying volumes for replication.
Choose the disks you want to replicate. Your disks are: c:
To replicate some of the disks, type the path of the disks, separated with a comma (for example, C:,D:). To replicate all disks, press Enter:
Disk to replicate identified: c:0 of size 30 GiB
```

The installer confirms that all of the disks were successfully identified.

A screenshot of a Windows command prompt window titled "C:\Users\Administrator\Downloads\AwsReplicationWindows\Installer.exe". The text inside the window shows the installation progress: "The installation of the AWS Replication Agent has started.", "AWS Region Name: eu-west-1", "AWS Access Key ID:", "AWS Secret Access Key:", "Verifying that the source server has enough free disk space to install the AWS Replication Agent. (a minimum of 2 GB of free disk space is required)", "Identifying volumes for replication.", "Choose the disks you want to replicate. Your disks are: c:", "To replicate some of the disks, type the path of the disks, separated with a comma (for example, C:,D:). To replicate all disks, press Enter:", "Disk to replicate identified: c:\0 of size 30 GiB", and "All volumes for replication were successfully identified.".

```
C:\Users\Administrator\Downloads\AwsReplicationWindows\Installer.exe
The installation of the AWS Replication Agent has started.
AWS Region Name: eu-west-1
AWS Access Key ID:
AWS Secret Access Key:
Verifying that the source server has enough free disk space to install the AWS Replication Agent.
(a minimum of 2 GB of free disk space is required)
Identifying volumes for replication.
Choose the disks you want to replicate. Your disks are: c:
To replicate some of the disks, type the path of the disks, separated with a comma (for example, C:,D:). To replicate all
disks, press Enter:
Disk to replicate identified: c:\0 of size 30 GiB
All volumes for replication were successfully identified.
```

Note

When identifying specific disks for replication, do not use apostrophes, brackets, or disk paths that do not exist. Type only existing disk paths. Each disk that you selected for replication is displayed with the caption **Disk to replicate identified**. However, the displayed list of identified disks for replication may differ from the data you entered. This difference can be due to several reasons:


- The root disk of the source server is always replicated, whether you select it or not. Therefore, it always appears on the list of identified disks for replication.
- AWS Transform MGN replicates whole disks. Therefore, if you choose to replicate a partition, its entire disk appears on the list and is later replicated. If several partitions on the same disk are selected, then the disk only appears once on the list.
- Incorrect disks may be chosen by accident. Ensure that the correct disks have been chosen.

⚠ Important

If disks are disconnected from a server, AWS Transform MGN can no longer replicate them, so they are removed from the list of replicated disks. When they are reconnected, the AWS Replication Agent cannot know that these were the same disks that were disconnected and therefore does not add them automatically. To add the disks after they are reconnected, rerun the AWS Replication Agent installer on the server.


Note that the returned disks need to be replicated from the beginning. Any disk size changes are automatically identified, but also cause a resync. Perform a test after installing the Agent to ensure that the correct disks have been added.

5. After all of the disks to be replicated have been successfully identified, the installer downloads and installs the AWS Replication Agent on the source server.



```
C:\Users\Administrator\Downloads\AwsReplicationWindows\Installer.exe
The installation of the AWS Replication Agent has started.
AWS Region Name: eu-west-1
AWS Access Key ID:
AWS Secret Access Key:
Verifying that the source server has enough free disk space to install the AWS Replication Agent.
(a minimum of 2 GB of free disk space is required)
Identifying volumes for replication.
Choose the disks you want to replicate. Your disks are: c:
To replicate some of the disks, type the path of the disks, separated with a comma (for example, C:,D:). To replicate all disks, press Enter:
Disk to replicate identified: c:0 of size 30 GiB
All volumes for replication were successfully identified.
Downloading the AWS Replication Agent onto the source server... Finished.
Installing the AWS Replication Agent onto the source server... Finished.
```

6. Once the AWS Replication Agent is installed, the server is added to the AWS Transform MGN console and undergoes the initial sync process. The installer provides the source server's ID.

A screenshot of a Windows command prompt window titled "C:\Users\Administrator\Downloads\AwsReplicationWindowsInstaller.exe". The window displays the following text: "The installation of the AWS Replication Agent has started. AWS Region Name: eu-west-1 AWS Access Key ID: AWS Secret Access Key: Verifying that the source server has enough free disk space to install the AWS Replication Agent. (a minimum of 2 GB of free disk space is required) Identifying volumes for replication. Choose the disks you want to replicate. Your disks are: c: To replicate some of the disks, type the path of the disks, separated with a comma (for example, C:,D:). To replicate all disks, press Enter: Disk to replicate identified: c:\0 of size 30 GiB All volumes for replication were successfully identified. Downloading the AWS Replication Agent onto the source server... Finished. Installing the AWS Replication Agent onto the source server... Finished. Syncing the source server with the Application Migration Service Console... Finished. The following is the source server ID: s- The AWS Replication Agent was successfully installed. Press Enter to close..."

```
C:\Users\Administrator\Downloads\AwsReplicationWindowsInstaller.exe
The installation of the AWS Replication Agent has started.
AWS Region Name: eu-west-1
AWS Access Key ID:
AWS Secret Access Key:
Verifying that the source server has enough free disk space to install the AWS Replication Agent.
(a minimum of 2 GB of free disk space is required)
Identifying volumes for replication.
Choose the disks you want to replicate. Your disks are: c:
To replicate some of the disks, type the path of the disks, separated with a comma (for example, C:,D:). To replicate all
disks, press Enter:
Disk to replicate identified: c:\0 of size 30 GiB
All volumes for replication were successfully identified.
Downloading the AWS Replication Agent onto the source server... Finished.
Installing the AWS Replication Agent onto the source server... Finished.
Syncing the source server with the Application Migration Service Console... Finished.
The following is the source server ID: s-
The AWS Replication Agent was successfully installed.
Press Enter to close...
```

You can review this process in real time on the **Source servers** page. [Learn more about the initial sync process.](#)

Installing the Agent on a secured network

The AWS Transform MGN AWS Replication Agent installer needs network access to MGN and Amazon S3 endpoints. If your on premises network is not open to MGN and Amazon S3 endpoints, then you can install the Agent with the aid of PrivateLink.

You can connect your on premises network to the subnet in your staging area VPC using AWS Virtual Private Network or AWS Direct Connect. To use the Site-to-Site VPN or Direct Connect, you must [use private IP in the replication settings](#) (not supported for IPv6).

Create a VPC endpoint for AWS Transform MGN

To allow the AWS Replication Agent installer to communicate with MGN, create an interface VPC endpoint for MGN in your staging area subnet. For more information, see [Creating an interface endpoint](#) in the *Amazon VPC User Guide*.

If the AWS replication agents are installed with a principal using [AWSApplicationMigrationAgentInstallationPolicy](#) and a VPCE policy is used (to scope down access), add this statement to your policy:

```
{
    "Effect": "Allow",
    "Principal": "*",
    "Action": "execute-api:Invoke",
    "Resource": "arn:aws:execute-api:<region>:*:*/POST/CreateSessionForMgn"
}
```

Use the created VPC Endpoint for AWS Transform MGN

Once you have created the VPC Endpoint, the AWS Replication Agent can connect to MGN via Site-to-Site VPN/Direct Connect by using the `--endpoint` installation parameter. Learn more about [Private DNS for interface endpoints](#) in the *Amazon VPC User Guide*.

Run the AWS Replication Agent installer with the `--endpoint` parameter. Enter your endpoint-specific DNS hostname within the parameter. The installer is then able to connect to MGN via the endpoint over your Site-to-Site VPN/Direct Connect connection.

Create an Amazon S3 endpoint for AWS Transform MGN

To allow the AWS Replication Agent installer to communicate with Amazon S3, create an interface Amazon S3 endpoint for MGN in your staging area subnet. For more information, see [Endpoints for Amazon S3](#) in the *Amazon VPC User Guide*.

Use the created Amazon S3 Endpoint for AWS Transform MGN

Once you have created the Amazon VPC Endpoint, the AWS Replication Agent can connect to Amazon S3 via Site-to-Site VPN/Direct Connect by using the `--s3-endpoint` installation parameter. Learn more about [Private DNS for interface endpoints](#) in the *Amazon VPC User Guide*.

Run the AWS Replication Agent installer with the `--s3-endpoint` parameter. Enter your endpoint-specific DNS hostname. The installer is then able to connect to Amazon S3 via the endpoint over your Site-to-Site VPN/Direct Connect connection.

Prescriptive guidance

A detailed guide for rehosting servers using MGN over private networks is available here:

[Migrating on-premises servers to AWS over private networks by using MGN.](#)

Uninstalling the Agent

Uninstalling the AWS Replication Agent from a source server stops the replication of that server. Uninstalling the AWS Replication Agent removes the source server from the AWS Transform MGN console.

Note

- The source server must be able to communicate with the MGN service in order for the Agent to be uninstalled successfully.
- If the Agent is uninstalled directly from a source server without disconnecting the server from MGN or finalizing the cutover within the MGN console, the replication metering period continues and once 2160 hours have elapsed, billing for replication begins.

Uninstalling the Agent through the AWS Transform MGN console

To uninstall the AWS Replication Agent through the AWS Transform MGN console:

Navigate to the **Source servers** page.

Check the box to the left of each server that you want to disconnect from MGN (by uninstalling the AWS Replication Agent). Open the **Actions** menu, and choose the **Disconnect from service** option to disconnect the selected server from MGN and AWS.

On the **Disconnect X server/s from service** dialog, choose **Disconnect**.

The AWS Replication Agent is uninstalled from all of the selected source servers. You can then archive these servers. [Learn more about archiving.](#)

Uninstalling the Agent manually through the source server

To uninstall the AWS Replication Agent manually through the source server:

Windows

Copy this folder to a new location: C:\Program Files (x86)\AWS Replication Agent\dist

From the new location, run in CMD as an administrator:

```
install_agent_windows.exe --remove
```

Linux

Run as root or with sudo these commands:

```
/var/lib/aws-replication-agent/uninstall-agent.sh
```

Reinstalling the Agent

To reinstall the AWS Replication Agent, download the latest version of the agent and follow the installation instructions. The AWS Replication Agent can be installed over an existing agent installation. You do not need to remove any previous versions prior to install.

- [Installing the AWS Replication Agent on Linux servers](#)
- [Installing the AWS Replication Agent on Windows servers](#)

Installing the AWS Transform MGN vCenter Client for Agentless Replication on vCenter source environments

AWS Transform MGN allows you to perform agentless snapshot replication from your vCenter source environment into AWS. This is achieved by installing the MGN vCenter Client in your vCenter environment. MGN recommends using agent-based replication when possible, as it supports CDP (Continuous Data Protection) and provides the shortest cutover window. Agentless replication should be used when your company's policies prevent you from installing the AWS Replication Agent on each individual server.

Topics

- [Agentless replication overview](#)
- [VMware limitations](#)
- [Generating vCenter Client IAM credentials](#)
- [Installing the MGN vCenter Client](#)
- [Replicating servers from vCenter to AWS](#)
- [Updating the vCenter or AWS Credentials](#)
- [Differentiating agentless and agent-based servers](#)

Agentless replication overview

Agentless snapshot-based replication allows you to replicate source servers on your vCenter environment into AWS without installing the AWS Replication Agent.

In order to use agentless replication, you must dedicate at least one VM in your vCenter environment to host the MGN vCenter Client. The MGN vCenter Client is a software bundle distributed by MGN and is available for installation as a binary installer. The installation process installs services on the client VM which allow MGN to remotely discover your VMs that are suitable for agentless replication, and to perform data replication between your vCenter environment and AWS through the use of periodic snapshot shipping.

Agentless snapshot based replication is divided into two main operations: discovery and replication:

The discovery process involves periodically scanning your vCenter environment to detect source server VMs that are suitable for agentless replication, and adding these VMs to the MGN console. Once a source server has been added, you may choose to initiate agentless replication on the source VM using the MGN API or console. The discovery process also collects all of the necessary information from vCenter in order to perform an agentless conversion process once a migration job is launched.

The replication process involves continuously starting and monitoring the “snapshot shipping processes” on the source server VM being replicated. A “snapshot shipping process” is a long running logical operation which consists of taking a VMware snapshot on the replicated VM, and launching an ephemeral replication agent process which uses VMware’s Changed Block Tracking (CBT) feature to identify changed volume data location, using Virtual Disk Development Kit (VDDK) to read the modified data, and sending the data from the source environment to the customer’s target AWS account. The first snapshot shipping process performs an “initial sync” which sends the entire disk contents of the replicating VM into AWS. Following snapshot shipping processes leverage CBT only to sync disk changes to the customer’s target AWS account. Each successful snapshot shipping process completes the replication operation by creating a group of consistent snapshots in the customer’s AWS account, which can then be used by the customer to launch test and cutover instances through the regular MGN mechanisms.

These are the main system components of agentless replication:


- **MGN vCenter Client** – A software bundle that is installed on a dedicated VM in your vCenter environment in order to facilitate agentless replication.

- vCenter Replication Agent – A java agent that is based on the AWS Replication Agent, which replicates a single VM using VDDK and CBT as the data source instead of the MGN driver (that is used by the AWS Replication Agent)
- MGN Service
- MGN console

Note

Agentless replication does not work in IPv6-only source environments.

This diagram illustrates the high level interaction between the different agentless replication system components:



Prerequisites

1. Ensure that you have initialized AWS Transform MGN before installing the MGN vCenter Client.

VMware limitations

- MGN supports using agentless and agent-based replication for migrations from VMC on AWS.
- MGN partially supports vMotion, Storage vMotion, and other features based on virtual machine migration (such as DRS and Storage DRS) subject to these limitations:
 - Migrating a virtual machine to a new ESXi host or datastore after one replication run ends, and before the next replication run begins, is supported as long as the vCenter account has sufficient permissions on the destination ESXi host, datastores, and datacenter, and on the virtual machine itself at the new location.
 - Migrating a virtual machine to a new ESXi host, datastore, and/or datacenter while a replication run is active, that is, while a virtual machine upload is in progress, is not supported. Cross vCenter vMotion is not supported for use with MGN.
- AWS does not provide support for migrating VMware Virtual Volumes.

Generating vCenter Client IAM credentials

In order to use the MGN vCenter Client, you must first generate the correct IAM credentials.

You need to create at least one AWS Identity and Access Management (IAM) user, and assign the proper permission policies to this user. Obtain an Access key ID and Secret access key, which you need to enter into the Agent installation prompt in order to begin the installation. We recommend that you use **IAM access last used information** to rotate and remove access keys safely. For more information, see [Rotating access keys](#).

1. Open the **AWS Management Console** and look for **IAM** under **Find Services**.
2. From the **IAM** main page, choose **Users** from the left-hand navigation menu.
3. You can either select an existing user or add a new user. To add a new user, choose **Add user**.
4. Give the user a **User name** and select the **Programmatic access** access type. Choose **Next: Permissions**.
5. Choose the **Attach existing policies directly** option. Search for **AWSApplicationMigrationVCenterClientPolicy** and **AWSApplicationMigrationAgentPolicy**. Select the policies and choose **Next: Tags**.
6. Add tags if you wish to use them and then choose **Next: Review**.
7. Review the information. Ensure that the **Programmatic access** type is selected and that the correct policy is attached to the user. Choose **Create user**.

8. A confirmation message appears and you can see the **Access key ID** and **Secret access key** that you need in order to install the AWS Replication Agent on your source servers.

To save this information as .csv file, choose **Download .csv**.

You can also access this information and re-generate your security credentials by navigating to **IAM > Users > Your user**.

Open the **Security credentials** tab and scroll down to **Access keys**. Here you can manage your access keys (create, delete, and more).

Installing the MGN vCenter Client

The first step to deploying the agentless solution is installing the MGN vCenter Client on your vCenter environment.

Note

If you have multiple vCenter environments, you need to install multiple clients. You may not have more than one MGN vCenter Client installed per AWS account. If you have multiple vCenter environments, you can either use a different AWS account for each environment or you can migrate your VMs serially, environment by environment, into the same AWS account.

After the MGN vCenter Client has been installed, it discovers all of the VMs in your vCenter environment and adds them to MGN.

MGN vCenter Client requirements

Ensure that you review the notes below prior to installing the MGN vCenter Client. Once you have read the notes, proceed to [install the client](#).

vCenter Client requirements

- You must install the MGN vCenter Client on a VM that has outbound network connectivity to the AWS Transform MGN API endpoints and outbound network connectivity to the vCenter endpoint. Customers who want to use PrivateLink can use VPN or AWS Direct Connect to connect to AWS.
- The MGN vCenter Client currently only supports VirtualDiskFlatVer2BackingInfo VMDK on CBT.

- You must log in to your Broadcom account and download VDDK 7.0.3.3 to the VM on which the MGN vCenter Client is installed. VDDK 7.0.3.3 must be used, regardless of the vCenter version used.
- The MGN vCenter Client requires these vCenter user permissions for agentless deployment. It is a best practice to create a dedicated role with these permissions and a dedicated user group with which the role is associated. Every new user created for the MGN vCenter Client needs to be a member of that group in order to obtain the required permissions. The vCenter predefined role: "Consolidated Backup user (sample) " provides most of these permissions. If that role is used, the **Toggle disk change tracking** permission must be provided..
 - Change configuration
 - Acquire disk lease
 - Toggle disk change tracking
 - Provisioning
 - Allow read-only disk access
 - Allow virtual machine download
 - Snapshot management
 - Create snapshot
 - Remove snapshot
- The VM on which the MGN vCenter Client is installed should meet these RAM, CPU, and memory requirements:
 - Minimal requirements (these requirements allow the replication of up to 5 servers in parallel) – 2 GiB RAM, 1 core, 10 GiB of free disk space
 - Optional performance requirements (these requirements allow the replication of the maximum number of 50 servers in parallel) – 16 GiB RAM, 8 cores, 10 GiB of free disk space
- VMs that are being replicated into AWS should have at least 2 GiB of free disk space.
- The VM on which the MGN vCenter Client is installed should not allow any incoming (ingress) traffic.
- The VM on which the MGN vCenter Client is installed should only allow outgoing traffic as following:
 - Egress TCP on the port on which the vCenter API runs.
 - Egress TCP on port 443 for communication with the MGN API.
 - Egress TCP on port 1500 – for the replication server.

- Patching of guest OS running the MGN vCenter Client should be handled by the customer as part of shared responsibility.
- IAM credentials used by the vCenter Client should be rotated on a regular schedule. Learn more about how to rotate access keys for IAM users in [this IAM blog post](#). IAM credentials can be regenerated by reinstalling the AWS Replication Agent.
- The VM that hosts the MGN vCenter Client should only be used for client hosting and should not be used for any other purposes.
- Only a trusted administrator should have access to the VM on which the MGN vCenter Client is installed.
- The MGN vCenter Client should be located in an isolated and dedicated network and considered a sensitive segment.
- You can deactivate the vCenter Client auto-update mechanism by running this command:
`touch /var/lib/aws-vcenter-client/.disable_auto_updates` Once auto-updates are deactivated, you need to reinstall the client to perform a manual update. If you deactivate the auto-update mechanism, you are responsible for ensuring that all security updates are performed on the client. After a manual update, you should validate the new hash against the [installer hash](#).

vCenter Client installer notes

- The MGN vCenter Client installer only supports vCenter 6.7, 7.0 and 8.0.
- The MGN vCenter Client can be installed on these 64 bit Linux versions:
 - Ubuntu 18.x+ (64 bit) - 22.04
 - Amazon Linux 2
 - RHEL 8.x
- If you are using a RHEL 8.x environment, ensure that you run the `sudo yum install python3` command to install python prior to launching the client installer.
- These flags are used by the installer:
 - usage: `aws-vcenter-client-installer-init.py [-h]`
 - `[--aws-access-key-id AWS_ACCESS_KEY_ID]`
 - `[--aws-secret-access-key AWS_SECRET_ACCESS_KEY]`
 - `[--region REGION]`
 - `[--endpoint ENDPOINT]`

- [--s3-endpoint S3_ENDPOINT]
- [--vcenter-host VCENTER_HOST]
- [--vcenter-port VCENTER_PORT]
- [--vcenter-user VCENTER_USER]
- [--vcenter-password VCENTER_PASSWORD]
- [--vcenter-ca-path VCENTER_CA_PATH]
- [--vddk-path VDDK_PATH]
- [--vcenter-client-tags KEY=VALUE [KEY=VALUE ...]]
- [--source-server-tags KEY=VALUE [KEY=VALUE ...]]
- [--disable-ssl-cert-validation]
- [--no-prompt]
- Use this flag for an unattended installation. If you are using this flag, you must also use the --force-delete-existing client flag.

[--force-delete-existing-client]

- Use this flag to delete an existing version of the vCenter Client from your VM. You must use this flag if you've previously installed the vCenter Client on the VM. If you use the --no-prompt flag, you must also use this flag.

[--version]

Optional arguments:

-h, --help show this help message and exit

vCenter environment requirements

- AWS Transform MGN supports VM hardware version 7 and higher with CBT activated. Ensure that you upgrade any VMs you have to hardware version 7 or higher. Ensure that CBT support is activated in your vSphere deployment. MGN activates CBT on replicating VMs. You can deactivate CBT after cutover.
- The VM being replicated into MGN must not contain any existing VMware snapshots.
- Once added to MGN, snapshot-based replication creates snapshots on the replicated VM, which may result in slower disk performance.

- VMs with independent disks, Raw Device Mappings (RDM), or direct-attach disks (iSCSI, NBD) are not supported for replication into MGN.
- The VM being replicated into MGN can be either stopped or running. Changing the VM state during data replication does not affect data replication and causes no data corruption.

MGN vCenter Client installation instructions

To install the MGN vCenter Client, follow these steps:

1. Download the MGN vCenter Client installer onto a VM within your vCenter environment. You can download the client from this URL: [https://aws-application-migration-service-\(region\).s3.\(region\).amazonaws.com/latest/vcenter-client/linux/aws-vcenter-client-installer-init.py](https://aws-application-migration-service-(region).s3.(region).amazonaws.com/latest/vcenter-client/linux/aws-vcenter-client-installer-init.py) Replace (region) with the AWS Region into which you are replicating.

This is an example of the installer link for us-east-1: <https://aws-application-migration-service-us-east-1.s3.us-east-1.amazonaws.com/latest/vcenter-client/linux/aws-vcenter-client-installer-init.py>

If you need to validate the installer hash, the correct hash can be found here: [https://aws-application-migration-service-hashes-\(region\).s3.\(region\).amazonaws.com/latest/vcenter-client/linux/aws-vcenter-client-installer-init.py.sha512](https://aws-application-migration-service-hashes-(region).s3.(region).amazonaws.com/latest/vcenter-client/linux/aws-vcenter-client-installer-init.py.sha512)

This is an example of the installer hash link for us-east-1: <https://aws-application-migration-service-hashes-us-east-1.s3.us-east-1.amazonaws.com/latest/vcenter-client/linux/aws-vcenter-client-installer-init.py.sha512>

2. In command prompt, navigate to the directory where you downloaded the MGN vCenter Client installer and run the installer with this command: `sudo python3 aws-vcenter-client-installer-init.py`
3. The installer prompts you for your credentials, enter the required info in each field and then press **Enter**:
 - AWS Access Key ID – Enter the AWS Access Key ID you generated in the previous section.
 - AWS Secret Access Key – Enter the AWS Secret Access Key you generated in the previous section.

- AWS Region name – The AWS Region of your account (for example, eu-west-1).
 - The Private Link endpoint for AWS Transform MGN (optional, leave blank if not using Private Link).
 - The VPC endpoint for Amazon S3 (optional, leave blank if not using a VPC endpoint).
4. The installer then prompts you to enter your vCenter information, enter the required info in each field and then press **Enter**:

- vCenter IP or hostname
- vCenter port (press Enter to use the default TCP Port 443)
- vCenter username
- vCenter password
- Path to vCenter root CA certificate (optional) - To use SSL certificate validation, download the certificates from <https://<vcenter-ip>/certs/download.zip> (example: `wget https://<vcenter-ip>/certs/download.zip --no-check-certificate`) then enter the path of the certificate (example: `/usr/local/src/lin/f7f2bd6e.0`). Otherwise, press **Enter** to deactivate SSL certificate validation.

 Note

- The certificate must be located in a file that's readable to the vCenter client user, such as a shared directory. If the certificate is not located in a shared directory, you see a permission error in the logs (Error 13).
- To use a certificate in your vCenter environment, you must setup a connection using a hostname. Using an IP does not work with a certificate.
- It's a security best practice to use certificates. Customers that do not use certificated authentication are responsible for any security issues that may arise.

- Path to VDDK tarball - Provide the path to the VDDK tarball that you previously downloaded onto the VM. (example: `path/to/VMware-vix-disklib-7.0.3-21933544.x86_64.tar.gz`). You can download VDDK tarball from your Broadcom account.
- Resource tags for the MGN vCenter Client (optional) - Use this format for tagging:

KEY=VALUE [KEY=VALUE ...] add resource tags to the MGN vCenter Client; use a space to separate each tag (e.g., --vcenter-client-tags tag1=val1 tag2=val2 tag3=val3)

- Resource tags for source servers to be discovered by the MGN vCenter Client (optional) - Use this format for tagging:

KEY=VALUE [KEY=VALUE ...] add resource tags to the source servers added by discovery; use a space to separate each tag (e.g., --vcenter-client-tags tag1=val1 tag2=val2 tag3=val3)

5. The installer downloads and installs the MGN vCenter Client and registers it with AWS Transform MGN.

6. Once the MGN vCenter Client has been installed, all of the VMs in your vCenter are added to AWS Transform MGN. The VMs are added in the DISCOVERED state.

 Note

- If you have a significant number of VMs in your vCenter environment, it may take some time for all of the VMs to become visible in the MGN console.
- The MGN vCenter Appliance is excluded from the discovered servers list.

You can configure transparent proxy either by using an environment variable prior to the installation (Linux and Windows), or by using the --proxy-address flag in the Linux installer:

- Using the installer: ./aws-vcenter-client-installer-init.py --proxy-address http://PROXY:PORT/
- Using environment variable: export https_proxy=http://PROXY:PORT/; ./aws-vcenter-client-installer-init.py

Make sure the proxy has a trailing forward slash.

Replicating servers from vCenter to AWS

Once you have successfully installed the MGN vCenter Client, all of your vCenter VMs are added to MGN in the DISCOVERED state. The DISCOVERED state means that the VM has not been replicated to AWS.

Note

VMware only sends data for up to 50 servers in parallel. Replicating more than 50 servers at once causes the rest to be queued and results in a longer wait.

By default, the MGN console only shows active servers. You can tell which servers are being shown by looking at the filtering box under the main **Source servers** title.

To see your discovered non-replicating servers that have been added from vCenter, open the filtering menu and choose **Discovered source servers**.

You now see all of your non-replicating DISCOVERED VMs.

To replicate one or more VMs into AWS, select the box to the left of each VM name, choose the **Replication** menu, and then choose **Start data replication**.

Choose **Start** on the **Start data replication for x servers** dialog.

The MGN console indicates that data replication has started.

To view the data replication progress, open the filtering menu and return to the default **Active source servers** view.

You now only see your replicating source servers. You can follow the launch process on the main **Source servers** view.

Once the VM has reached the **Ready for testing** state under **Migration lifecycle**, you can continue to [launch test and cutover instances](#) and perform all other regular MGN operations on the server.

Updating the vCenter or AWS Credentials

Users who want to change the vCenter or AWS credentials used by the MGN appliance should follow these steps. This change requires root privileges on the appliance:

1. In the command prompt, navigate to the aws-vcenter-client directory:

```
cd /var/lib/aws-vcenter-client/1.1.8/
```

2. Run the vCenter configuration update tool with this command:

```
sudo ./vcenter_configuration_update
```

3. When running the vCenter configuration update tool, you are prompted to provide the necessary credentials. Follow these steps to update the credentials. Provide the required info in each field and then press Enter:
 - New vCenter username (`--new-vcenter-username`)
 - New vCenter password (`--new-vcenter-password`)
 - New AWS Access Key ID (`--new-aws-access-key-id`)
 - New AWS Secret Access Key (`--new-aws-secret-access-key`)
 - New path to the CA (optional) (`--new-ca-path`)
4. If you do not provide the `--new-ca-path`` flag, the tool first asks if you want to update the CA path. If you answer yes, it prompts you for the new CA path. If you answer no, the tool uses the CA path from the previous configuration. The tool verifies the new vCenter and AWS credentials by attempting to connect to vCenter and MGN using them.
5. Upon successful connection to vCenter and MGN, the tool saves the new credentials and restarts the necessary services.
6. In case of failure to connect to vCenter or MGN, the new credentials are not stored, and the previous configuration is retained. This error message is displayed: Failed to connect to the vCenter endpoint or MGN using the new connection details. The configuration changes will not be applied.

Differentiating agentless and agent-based servers

You can differentiate an agentless vCenter VM that's replicating through snapshot shipping and an agent-based server (from any source infrastructure) through several ways:

1. On the **Source servers** page, under the **Replication type** column, the MGN console identifies the replication type, whether it is through **Snapshot shipping** (agentless) or **Agent based**.
2. In the server details view, under the **Migration dashboard**, agentless servers that are replicated through snapshot shipping have an additional **Lifecycle** step – **Not started**.
3. Similarly, in the server details view, under the **Migration dashboard**, the **Data replication status** box shows the **Replication type** as **Snapshot shipping**.

Manage source servers

The **Source servers** page lists all of the source servers that have been added to AWS Transform MGN (AWS MGN). The **Source servers** page allows you to manage your source servers and perform a variety of commands for one or more servers (such as controlling replication and launching test and cutover instances). The **Source servers** page is the main page of AWS MGN console and you will most likely interact with AWS Transform MGN predominantly through this page.

Interacting with the source servers page

The **Source servers** page shows a list of source servers. Each row on the list represents a single server. The table displays **Active source servers** by default. You can use the dropdown menu above the first row to choose to display:

- **Agentless discovered servers** - servers discovered by the agentless replication process, which are not yet replicating.
- **Imported servers** - servers whose information was imported into MGN, but an agent is not yet installed on them.
- **Archived source servers** - servers whose migration has completed, and the users archived their data.

The **Source servers** page provides key information for each source server under each of the columns on the page.

The columns include:

- **Selector column** – This blank checkbox selector column allows you to select one or more source servers. When a server is selected, you can interact with the server through the **Actions**, **Replication**, and **Test and cutover** menus. Selected servers are highlighted.
- **Source server name** – This column shows the unique server name for each source server.
- **Alerts** – This column shows whether any alerts exist for the server.
 - **No indication** – a healthy server for which a test or cutover instance has not been launched.
 - **Launched** – a healthy server for which a test or cutover instance has been launched.
 - **A clock icon with a warning message** – a server that is experiencing a temporary issue such as lag or backlog
 - **A red x and message** – a server that is experiencing significant issues, such as a stall.

- **Replication type** – This column identifies whether the server is being replicated through the default **Agent based** replication or through **Snapshot shipping**. [Learn more about agentless based snapshot shipping replication for vCenter](#).
- **Migration lifecycle** – This column shows the migration lifecycle state for each source server. This way you can easily know which lifecycle step the server is undergoing. Migration lifecycle steps include these: [Learn more about Migration lifecycle steps](#).
 - **Not ready**
 - **Ready for testing**
 - **Test in progress**
 - **Ready for cutover**
 - **Cutover in progress**
 - **Cutover complete**
 - **Disconnected**
- **Data replication status** – This column shows the current status of data replication for the server. The information presented in this column changes based on the server's Migration lifecycle state and whether the server is experiencing any issues.

This column shows a variety of information, including:

- **Not Started** – Data replication has not started.
- **Paused** – Data replication has been paused.
- **Healthy** – The server is healthy and ready for Test or Cutover instance launch.
- **Done** – The server has been successfully cutover and data replication has been stopped as a result.
- **Percentage complete** – The percentage of the server's storage that was successfully replicated if the server is undergoing initial sync or a rescan.
- **Lag** – Whether server is experiencing any lag. If it is; the lag time is indicated.
- **Backlog** – Whether there is any backlog on the server (in MiB).
- **Last snapshot** – This column shows the time the last consistent snapshot was taken of the source server. Servers that are still in the initial sync process and those that have been disconnected do not show any info in this field. Healthy servers show a recent snapshot. Unhealthy servers' last snapshot indicates the last time they were healthy.

- **Next step** – This column shows the next step that needs to be undertaken in order to successfully complete a cutover for the server. The information presented in this column changes based on the server's Migration lifecycle state and whether the server is experiencing any issues.

This column shows a variety of next steps, including:

- **Wait for initial sync to complete** – Data replication is initiating for the server. Wait for the initial sync process to complete.
- **Start data replication** – Data replication has not been started or is paused.
- **Launch test instance** – The server is ready to launch a test instance.
- **Complete testing and mark as 'ready for cutover'** – The server has launched a Test instance that needs to be reverted or finalized.
- **Launch cutover instance** – The server had a Test instance launched and finalized and now is ready to launch a cutover instance.
- **Finalize cutover** – The server has launched a cutover instance that needs to be finalized.
- **Terminate launched instance, Launch cutover instance** – The server is ready for cutover. Terminate the launched Test instance and launch a cutover instance.
- **Resolve cause of stalled data replication** – The server is experiencing significant issues such as a stall that need to be addressed.
- **Wait for lag to disappear, then X.** – The server is experiencing a temporary lag. Wait for the lag to disappear and then perform the indicated action.
- **Mark as archived** – The server has been successfully cut over and can now be archived.

Topics

- [Add a source server](#)
- [Manage data replication](#)
- [Manage test and cutover instances](#)
- [Source server migration metrics](#)
- [Filtering the source servers list](#)

Add a source server

To add a server, simply choose **Add server**.

The **Add servers** prompt opens, allowing you to create a custom installation command by taking these steps:

- 1. Select your operating system. The installation command is different for Windows and Linux

 Note

If you want to install the AWS Replication Agent on a legacy Windows OS (Windows Server 2003, Windows Server 2008 or Windows Server 2008 R2), you must choose the **Legacy OS: Windows Server 2003 or Windows Server 2008** box. This downloads a unique version of the AWS Replication Agent installer that is only valid for legacy Windows OSs (`AwsReplicationWindowsLegacyInstaller.exe`). **Do not** use this installer file to install the agent on any other OS types.

- 2. Select your replication preferences for the source server. The selected preferences are added as installation prompts to the custom installation command that is generated by this form.

Choose the **Replicate all disks** option to replicate all of the disks of the source server. This is the default option. This option adds the `--no-prompt` prompt to the installation command.

Select the **Choose which disks to replicate** option to choose which specific disks you want to replicate. You are prompted to select which disks to replicate during agent installation.

- 3. Enter the credentials [you previously generated for AWS Replication Agent installation](#). The form does not send the secret, but does add it to the installation command.
- 4. If you are adding a Windows source server to AWS MGN, download the installer onto the source server. The installer is downloaded from the AWS Region of your account. If you're adding a Linux source server, skip this step.
- 5. Copy the generated custom installation command and input it into the command line on your source server. Proceed with [AWS Replication Agent installation as instructed in the documentation](#).

To run a variety of commands on your source servers, select one or more servers and choose the **Actions, Replication, or Test and cutover** menu.

Edit server

To edit a source server, select the server and choose **Edit**. You can edit the following fields:

- **User-provided ID** – Edit the user-provided identifier for the source server. This value is case insensitive.
- **FQDN** – Edit the fully qualified domain name for the source server.
- **Platform** – Edit the platform of the source server. This field is editable only for imported servers. Once a server begins replicating, the platform is automatically detected and cannot be modified.

Actions menu

The **Actions** menu allows you to perform these actions:

- **View server details** – Choose this option to see the server details view for the selected server. This option is only available when a single server is selected.
- **Add server to application** – Choose this option to easily add servers to an application.
- **Disconnect from service** – Choose this option to disconnect the selected server from MGN and AWS. This option should be used when data replication is complete.

On the **Disconnect X server/s from service** dialog, choose **Disconnect**.

Important

This uninstalls the AWS Replication Agent from the source server and data replication stops for the source server. If you need to restart data replication for this server, you need to reinstall the agent. This action does not affect any test or cutover instances that have been launched for this source server, but you can no longer identify which source servers your Amazon EC2 instances correspond to.

- **Mark as archived** – Choose this option in order to archive the server. You should only archive servers for which you have already performed a cutover. Archived servers are removed from the main **Source servers** page, but can still be accessed through filtering options.

On the **Archive X servers** dialog, select **Archive**.

To see your archived servers, open the **Preferences** menu by choosing the gear button.

Select the **Show only archived servers** option and choose **Confirm**.

You are now able to see all of your archived servers. Unselect this option to see your non-archived servers.

- **Delete from service** – Choose this option to remove the source server from the database. This option supports single or multiple server selection. Only servers that are not actively replicating can be deleted.

Manage data replication

You can manage data replication for the source server through these actions on the **Replication** drop-down menu:

- **Start data replication** – Restarts data replication for a source server on which data replication has been stopped. If you are using agent-based replication you don't necessarily have to reinstall the agent.
- **Stop data replication** – Stops data replication for a source server. Stopping the replication stops billing, deletes existing snapshots and storage volumes (Amazon EBS volumes or FSx for ONTAP volumes), and terminates replication servers. Configuration is retained, and if you are using agent-based replication the agent is not uninstalled.
- **Pause data replication** – Pauses data replication for a source server. Pausing the replication does not stop billing or delete existing snapshots or storage volumes. Replication servers are not terminated and if you are using agent-based replication the agent is not uninstalled from the source server.
- **Resume data replication** – Resume data replication for a paused source server. This syncs any changes since the last synchronization and completes the data replication flow.

Choose **Edit** in the **Replication settings** section to access the **Edit Replication Settings** page, where you can edit the settings for the selected source server. [Learn more about editing replication settings.](#)

Manage test and cutover instances

The **Test and cutover** menu allows you to manage your test and cutover instances. For a more in-depth step-by-step guide to launching test and cutover instances, see the [Launching test and cutover instances documentation](#).

- **Launch test instances** – Choose this option to launch a test instance for this server.

When the **Launch test instances for X servers** dialog appears, choose **Launch** to begin the test.

The AWS Transform MGN console indicates **1 launch job complete** after the test has been completed successfully.

- **Mark as "Ready for cutover"** – Use this option to finalize testing for this server after you have completed all of the necessary tests in preparation for cutover.

When the **Mark X servers as "Ready for cutover"** dialog appears, select whether you want to terminate the launched instances used for testing. We recommend that you terminate these instances, as you will be charged for them even though you no longer need them. Check the **Yes, terminate launched instances (recommended)** box and choose **Continue**.

The AWS Transform MGN console indicates that testing has been finalized. The selected source servers' **Migration lifecycle** column shows the **Ready for cutover** status and the launched Test instances are deleted if that option was selected.

- **Revert to "ready for testing"** – Choose this option to revert a finalized test for this server if you want to run further tests prior to initiating a cutover.

The **Revert testing for X servers** dialog appears. Select whether you want to terminate the launched instances used for testing. We recommend that you terminate these instances, as you will be charged for them even though you no longer need them. Check the **Yes, terminate launched instances (recommended)** box and choose **Revert**.

The AWS Transform MGN console indicates that testing has been reverted. The selected source servers' **Migration lifecycle** column shows the **Ready for testing** status and the launched Test instances are deleted if that option was selected.

- **Launch cutover instances** – Choose this option to launch a cutover instance for this server after you have finalized all of your testing and are ready to initiate a cutover.

The **Launch cutover instances for X servers** dialog appears. Choose **Launch** to begin the cutover.

The AWS Transform MGN console indicates **1 launch job complete** after the cutover has been completed successfully.

This changes your source servers' **Migration lifecycle** status to **Cutover in progress**, indicating that the cutover is in progress but has not yet been finalized.

- **Finalize cutover** – Choose this option to finalize the cutover for this server after you have successfully performed a cutover.

This changes your source servers' **Migration lifecycle** status to **Cutover complete**, indicating that the cutover is complete and that the migration has been performed successfully. In addition, this stops data replication and causes all replicated data to be discarded. All AWS resources used for data replication are terminated.

The **Finalize cutover for X servers** dialog appears. Choose **Finalize**.

The AWS Transform MGN console indicates **X servers cutover. Data replication has been stopped for servers** once the cutover has been completed successfully. The AWS Transform MGN console automatically stops data replication for the cutover source servers to save resource costs. The selected source servers' **Migration lifecycle** column shows the **Cutover complete** status, the **Data replication** column shows **Disconnected** and the **Next step** column states **Mark as archived**. The source servers have now been successfully migrated into AWS and can be archived.

 Note

This action does not uninstall the AWS Replication Agent from the source server. Use the **Disconnect from service** option under the **Actions** menu when you have completed the migration and want to uninstall the agent from your source server.

- **Revert to "ready for cutover"** – Choose this option to revert a finalized cutover for this server if you encounter any issues or want to reverse the cutover for any reason.

This reverts your source servers' **Migration lifecycle** to the **Ready for cutover** status, indicating that these servers have not undergone cutover.

The **Revert cutover for X servers** dialog appears. Choose **Revert**.

- **Edit launch settings** – Use this option to edit the launch settings for this server. You are redirected to the **Edit launch settings** page. [Learn more about launch settings.](#)
- **Edit post-launch settings** – Use this option to edit the post-launch settings for the selected source server or group of source servers. [Learn more about post-launch settings.](#)
- **Terminate launched instance** – Choose this option if you want to delete your test or cutover instance for any reason at any time. It can only be selected for a server that has a launched test or cutover instance.

When the **Terminate launched instance** dialog appears, choose **Terminate**.

Source server migration metrics

The source server migration metrics present an aggregated overview of your source servers, focused on three topics: **Alerts**, **Data replication status**, and **Migration lifecycle**.

Understand source server alerts

The source server **Alerts** migration metric presents an aggregated overview of the application associated servers alerts. You can look up an individual source server's **Alerts** in the **Source servers** table at the bottom of the page.

- A healthy server for which a test or cutover instance has not been launched displays a **Healthy** status.
- A healthy server for which a test or cutover instance has been launched displays a **Healthy** status.
- A server that is experiencing a temporary issue such as lag or backlog displays a **Lagging** status.
- A server that is experiencing significant issues, such as a stall, displays a **Stalled** status.

Understand data replication status

The source server **Data replication status** migration metric presents an aggregated overview of your source servers' data replication status. You can look up an individual source server's **Data replication status** status in the **Source servers** table at the bottom of the page.

Source server **Data replication status** can have one of these values:

- **Transferring snapshot**
- **Initial sync**
- **Finalizing sync**
- **Lagging**
- **Healthy**
- **Stalled**
- **Rescanning**
- **Not started**
- **Initiating**

- **Creating snapshot**
- **Paused**
- **Disconnected**

Understand the migration lifecycle metric

The source server **Migration lifecycle** metric shows an aggregated overview of your source servers' migration lifecycle. You can look up an individual source server's **Migration lifecycle** status in the **Source servers** table at the bottom of the page.

Source server **Migration lifecycle** can have one of these values:

- **Stopped**
- **Not ready**
- **Ready for testing**
- **Test in progress**
- **Ready for cutover**
- **Cutover in progress**
- **Cutover complete**
- **Disconnected**
- **Discovered**

Filtering the source servers list

You can customize the **Source servers** page through filtering. Filtering allows you to easily filter your servers by one or multiple properties.

Choose the **Filter servers** field and choose the filtering property from the **Properties** menu.

You can filter by a variety of properties, including:

- Alerts – Filter by specific alert (lagging, stalled, launched).
- Hostname – Filter by a specific hostname or a specific string of characters.
- Migration lifecycle – Filter by the migration lifecycle state.
- Data replication status – Filter by the data replication status.

- Next step – Filter by next step.
- SourceServerID – Filter by specific source server ID or string.
- Tags – Filter by tags. Relevant specific tag values appear under the **Tags** category. Choose the value to filter by.

You can filter by multiple properties at once in order to narrow down your results.

Choose **Clear filters** to clear the current filtering properties selected.

Access details on a source server

You can access the server details view by choosing the **Source server name** of any server on the **Source servers** page.

You can also access the server details view by checking the box to the left of any single source server on the **Source servers** page and choosing **Actions > View server details**.

The server details view shows information and options for an individual server. Here, you can fully control and monitor the individual server.

You can also perform a variety of actions, control replication, and launch test and cutover instances for the individual server from the server details view.

The **Next actions** box serves as a helpful guide to the state of the server and the next steps you need to take in order to complete the migration process for the server.

The server details view is divided into several tabs, including:

- Migration dashboard
- Server info
- Tags
- Disk settings
- Replication settings
- Launch settings

Topics

- [Monitor the server in the migration lifecycle](#)

- [Review source server information](#)
- [Add or remove tags assigned to source servers](#)
- [Review disk settings for source servers](#)
- [Change target storage type](#)
- [Edit replication settings for a source server](#)
- [Review launch settings for a source server](#)
- [Review post-launch settings](#)
- [Editing the post-launch settings](#)
- [Activating and deactivating post-launch actions](#)
- [Deploying post-launch actions](#)

Monitor the server in the migration lifecycle

The **Migration dashboard** tab allows you to monitor the server in relation to the migration lifecycle.

Here, you can see the lifecycle state the source server is currently on, a detailed view of the data replication status, and any events that the source server has undergone (in AWS CloudTrail). You can use the Migration dashboard to monitor the status of your source server and to troubleshoot migration and data replication issues.

Understand lifecycle states

The **Lifecycle** view shows the current state of each server within the migration lifecycle.

Lifecycle states include:

- **Not ready** – The server is undergoing the Initial Sync process and is not yet ready for testing. Data replication can only commence once all of the Initial Sync steps have been completed.
- **Ready for testing** – The server has been successfully added to AWS Transform MGN and data replication has started. Test or cutover instances can now be launched for this server.
- **Test in progress** – A test instance is currently being launched for this server.
- **Ready for cutover** – This server has been tested and is now ready for a cutover instance to be launched.

- **Cutover in progress** – A cutover instance is currently being launched for this server.
- **Cutover complete** – This server has been cutover. All of the data on this server has been migrated to the cutover instance.

The lifecycle always displays the **Launch status**, **Last test**, and **Cutover status** for each server that has undergone these stages.

Topics

- [Not ready](#)
- [Ready for testing](#)
- [Test in progress](#)
- [Ready for cutover](#)
- [Cutover in progress](#)
- [Cutover complete](#)

Not ready

The **Not ready** lifecycle state represents several possible scenarios:

Topics

- [Server undergoing initial sync](#)
- [Unable to complete initiation](#)

Server undergoing initial sync

A source server that has been added to AWS Transform MGN automatically begins the initial sync process after AWS Replication Agent installation.

Data replication can only commence after all of the initial sync steps have been completed. The server is in the **Not ready** lifecycle state until initial sync has been successfully completed.

Initial sync steps include:

- Initiation
 - Creating firewall rules
 - Creating replication server

- Booting replication server
- Resolving Service Manager address
- Authenticating with the Service Manager
- Downloading replication software
- Creating staging disks
- Pairing replication server with agent
- Establishing communication between AWS Replication Agent and replication server
- Sync (0% to 100%)
- Flush backlog (if any)
- Create first launchable snapshot

You can review the overall progress of the Initial Sync process under the **Data replication status** view.

It provides the percentage of completion, the time left until initial sync is finished, and whether there are any issues (such as a stall).

You can tell that a server has successfully completed the initial sync process through several indicators on the main **Source servers** page as well as on the **Migration dashboard** tab for an individual server.

On the main **Source servers** page, a newly added server that has completed initial sync for the first time shows **Ready for testing** under the **Migration lifecycle** column and **Healthy** under the **Data replication status** column.

On the individual server view, under the **Migration dashboard** tab, the **Lifecycle** section shows the **Ready for testing** status. The **Data replication status** section shows the **Healthy** status.

 Note

Servers automatically undergo initial sync every time there is a network disconnect.

Unable to complete initiation

The server is in the **Not ready** Lifecycle state until Initial Sync has been successfully completed.

If the Initial Sync process is stalled for any reason, the **Data replication status** section indicates that replication has stalled.

Scroll down to the Replication initiation steps to see the step on which the error occurred. The step on which initial sync failed is marked with a red "x".

You must fix the indicated issue before the initial sync process can continue. You are not able to migrate your server and the server remains in the **Not ready** state until the issue is resolved.

Each step has troubleshooting methods.

Ready for testing

Once the server has successfully completed the Initial Sync process, it enters the **Ready for testing** lifecycle state.

The **Data replication status** box shows a **Healthy** state, indicating that the server is healthy.

You can now launch a test instance for this server. The server stays in the **Ready for testing** lifecycle state until you launch a Test instance for the server.

Test in progress

Once you have launched a Test instance for your server, the migration dashboard shows the **Test in progress** lifecycle state.

Within the **Lifecycle** box, you can review the **Launch status** and **Last test** information fields for the test instance.

- The **Launch status** field shows the time of the test instance launch. While the Testing instance is being launched, the **Launch status** field shows **Waiting**.
- Once the test instance has been launched, the Launch status shows **Launched**. Wait for the instance to boot and then choose **View in EC2 console** link to open the EC2 console in a new tab, in order to view and monitor your launched test instance.
- The AWS EC2 console opens in a new tab and automatically searches for and displays your test instance.
- The **Last test** field shows the date of the last test. To review the test launch details, choose **Job ID**, which opens the job within the **Launch History** page in a new tab.
- On the main **Source servers** page, the **Migration lifecycle** column shows **Test in progress** and the **Next step** column shows **Complete testing and mark as "Ready for cutover"**.

- The server stays in the **Test in progress** Lifecycle state until you finalize your testing and mark the server as **Ready for cutover**.

You can use these indicators to verify that your test instance was successfully launched:

- On the **Server Details > Lifecycle** pane, the Launch status states **Launched**.
- On the main **Source servers** page, the **Alerts** column shows the **Launched** status.

Ready for cutover

After you have finalized your testing, the Migration dashboard shows the **Ready for cutover** lifecycle state.

- The **Launch status** field shows the time of the last test instance launch. Choose the **View in EC2 console** link to open the EC2 console in a new tab in order to view and monitor your launched Test instance.
- The **Last test** field shows the date the last test was started. You can review the test launch details by choosing the **Job ID**. This opens the relevant Job.
- The **Cutover** field shows the date of the last cutover instance launch, if applicable. You can review the cutover launch details by choosing the **Job ID**. This opens the relevant Job.
- On the **Source servers** page, the **Migration lifecycle** column shows **Ready for cutover** and the **Next step** column shows **Terminate test instance; Launch cutover instance**.

The server stays in the **Ready for cutover** Lifecycle state until you launch a cutover instance.

Cutover in progress

Once you have launched a cutover instance for your server, the Migration dashboard shows the **Cutover in progress** Lifecycle state.

- The **Launch status** field shows the last time of cutover launch. Choose the **View in EC2 console** link to open the EC2 console in a new tab in order to view and monitor your launched cutover instance.
- The **Last test** field shows the date the last test was started. You can review the test launch details by choosing the **Job ID**. This opens the Job.
- The **Cutover** field shows the date of the last cutover instance launch. You can review the cutover launch details by choosing the **Job ID**. This opens the Job.

- On the **Source servers** page, the **Migration lifecycle** column shows **Cutover in progress** and the **Next step** column shows **Complete the cutover**.

The server stays in the **Cutover in progress** Lifecycle state until you complete the cutover.

Cutover complete

Once you have completed your cutover instance launch for your server, the Migration dashboard shows the **Cutover complete** lifecycle state. This is the final state in the migration lifecycle. This state indicates that you have successfully migrated your source server to AWS.

- The **Launch status** field shows **Launched**. Choose the **View in EC2 console** link to open the EC2 console in a new tab in order to view and monitor your launched cutover instance.
- The **Last test** field shows the date the last test was started. You can review the test launch details by choosing the **Job ID**. This opens the Job.
- The **Cutover** field shows the date you finalized your Cutover instance launch. You can review the cutover launch details by choosing the **Job ID**. This opens the Job.
- The AWS Transform MGN console automatically stops data replication for the source servers that were cutover in order to save resource costs.
- On the **Source servers** page, the selected source servers' **Migration lifecycle** column shows the **Cutover complete** status, the **Data replication status** column shows **Disconnected** and the **Next step** column shows **Mark as archived**.

The lifecycle also shows the status of any post-launch actions for the server. [Learn more about post-launch actions](#).

Understand data replication states

The **Data replication status** section provides an overview of the overall source server status, including:

- **Replication progress** – The percentage of the server's storage that was successfully replicated.
- **Rescan progress** – The percentage of the server's storage that was rescanned (in the event of a rescan)
- **Total replicated storage** – The total amount of storage replicated (in GiB).
- **Lag** – Whether the server is experiencing any lag. If it is - the lag time is indicated.

- **Backlog** – Whether there is any backlog on the server (in MiB)
- **Elapsed replication time** – Time elapsed since replication first began on the server.
- **Last seen** – The last time the server successfully connected to AWS Transform MGN.
- **Replication start time** – The date and time replication first began on the server.

Data replication can be in one of several states, as indicated in the panel title:

- **Initial sync:** initial copying of data from external servers is not done. Progress bar and **Total replicated storage** fields indicate how far along the process is.
- **Healthy:** all data has been copied and any changes at the source are continuously being replicated (data is flowing).
- **Rescan:** an event happened that forced the agent on the external server to rescan all of the blocks on all of the replicated disks. This is similar to the initial sync but faster because only changed blocks are copied. A rescan progress bar is displayed.
- **Stalled:** data is not flowing. You may need to intervene. When stalled, either the initial sync never completes, or the difference between the state of the replicated data and that of the source server continues to grow. When the state is stalled, then the replication initiation checklist is shown and indicates where the error occurred that caused the stall.

This panel also shows:

- **Total replicated storage:** size of all disks being replicated for this source server, and how much has been copied to AWS (once initial sync is complete)
- **Lag:** if you launch a recovery instance now, how far behind it will be from the state at the source. Normally this should be none.
- **Backlog:** how much data has been written at source but has not yet been copied to AWS. Normally this should be none.
- **Last seen:** when is the last time the AWS Replication Agent communicated with the AWS Transform MGN service or the replication server.

If everything is working as it should and replication has finished initializing, the Data replication progress section shows a **Healthy** status.

If there are initialization, replication, or connectivity errors, the **Data replication status** section shows the cause of the issue (for example, a stall).

If the error occurred during the initialization process, then the exact step during which the error occurred is marked with a red "x" under **Replication initiation steps**.

Understand the state of post-launch actions

The **Post-launch actions** view shows the current execution status of post-launch actions.

The status includes:

- Name – the name of the action is a link to the detailed execution status in the [AWS Systems Manager](#) console.
- Execution status – provides the current action status.
- Start time – the start time of the execution of the action script. This column is empty for actions that have not yet started execution.
- End time – the end time of the execution of the action script. This column is empty for actions that have not yet completed execution.
- Details – error messages are shown in this column.
- Link – used by the “DR after migration” action. Provides a link to the replicated server in the [AWS Elastic Disaster Recovery](#) console.

Review events and metrics in AWS CloudTrail

You can review AWS Transform MGN events and metrics in AWS CloudTrail. Choose **View CloudTrail Event History** to open AWS CloudTrail in a new tab.

Learn more about [monitoring AWS MGN](#).

Learn more about AWS CloudTrail events in the [AWS CloudTrail user guide](#).

Understand server actions and replication control

You can perform a variety of actions, control data replication, and manage your testing and cutover for an individual server from the server details view.

Topics

- [Actions menu](#)
- [Replication menu](#)

- [Test and cutover menu](#)
- [Alerts and errors](#)

Actions menu

The **Actions** menu allows you to perform these actions:

- **Add servers** – Choosing the **Add servers** option opens the Add servers prompt, through which you can construct a custom installation command to use when adding Linux or Windows source servers.

To construct a custom installation command, take these steps:

1. Select your operating system. The installation command is different for Windows and Linux.

Note

If you want to install the AWS Replication Agent on a legacy Windows OS (Windows Server 2003, Windows Server 2008 or Windows Server 2008 R2), you must choose the **Legacy OS: Windows Server 2003 or Windows Server 2008** box. This downloads a unique version of the AWS Replication Agent installer that is only valid for legacy Windows OSs (AwsReplicationWindowsLegacyInstaller.exe). **Do not** use this installer file to install the agent on any other OS types.

2. Select your replication preferences for the source server. The selected preferences are added as installation prompts to the custom installation command that are generated by this form.

Choose the **Replicate all disks** option to replicate all of the disks of the source server. This is the default option. This option adds the `--no-prompt` prompt to the installation command.

Select the **Choose which disks to replicate** option to choose which specific disks you want to replicate. You are prompted to select which disks to replicate during agent installation.

3. Enter the credentials that [you previously generated for AWS Replication Agent installation](#). The form does not send the secret, but does add it to the installation command.
4. If you have not yet obtained the necessary credentials, [follow these instructions](#).
5. If you are adding a Windows source server to AWS MGN, download the installer onto the source server. The installer is downloaded from the AWS Region of your account. If you're adding a Linux source server, skip this step.

6. Copy the generated custom installation command and input it into the command line on your source server. Proceed with [AWS Replication Agent installation as instructed in the documentation](#).
- **View server details** – Choose this option to open the server details view for the selected server. This option is only available when a single server is selected.
- **Disconnect from service** – Choose this option to disconnect the selected server from MGN and AWS. This option disconnects the source server and should be used when data replication is complete.

On the **Disconnect X server/s from service** dialog, choose **Disconnect**.

Important

This uninstalls the AWS Replication Agent from the source server and data replication stops for the source server. If you need to restart data replication for this server, you need to reinstall the agent. This action does not affect any test or cutover instances that have been launched for this source server, but you are no longer able to identify which source servers your Amazon EC2 instances correspond to.

- **Mark as archived** – Choose this option to archive the server. You should only archive servers for which you have already performed a cutover. Archived servers are removed from the main **Source servers** page, but can still be accessed through filtering options.

On the **Archive X servers** dialog, select **Archive**.

To see your archived servers, open the **Preferences** menu by choosing the gear button. Select the **Show only archived servers** option and choose **Confirm**. You can now see all of your archived servers. Unselect this option to see your non-archived servers.

Replication menu

The Replication menu allows you to manage data replication for your source servers through these actions:

- **Edit replication settings** – Choose this option to be redirected to the **Edit replication settings** page, where you can edit specific replication settings for the selected source server. [Learn more about editing replication settings](#).

Test and cutover menu

The **Test and cutover menu** allows you to manage your test and cutover instances.

- **Launch test instances** – Choose this option to launch a test instance for this server.

When the **Launch test instances for X servers** dialog appears, choose **Launch** to begin the test.

The AWS Transform MGN console indicates **1 launch job complete** after the test has been completed successfully.

- **Finalize testing** – Choose the **Mark as "Ready for cutover"** option to finalize testing for this server after you have completed all of the necessary tests in preparation for cutover.

When the **Mark X servers as "Ready for cutover"** dialog appears, select whether you want to terminate the launched instances used for testing. We recommend that you terminate these instances, as you will be charged for them even though you no longer need them. Check the **Yes, terminate launched instances (recommended)** box and choose **Continue**.

The AWS Transform MGN console indicates that testing has been finalized. The selected source servers' **Migration lifecycle** column shows the **Ready for cutover** status and the launched test instances are deleted if that option was selected.

- **Revert to "ready for testing"** – Choose this option to revert a finalized test for this server if you want to run further tests prior to initiating a cutover.

When the **Revert testing for X servers** dialog appears, select whether you want to terminate the launched instances used for testing. We recommend that you terminate these instances, as you will be charged for them even though you no longer need them. Check the **Yes, terminate launched instances (recommended)** box and choose **Revert**.

The AWS Transform MGN console indicates that testing has been reverted. The selected source servers' **Migration lifecycle** column shows the **Ready for testing** status and the launched test instances are deleted if that option was selected.

- **Launch cutover instances** – Choose this option to launch a cutover instance for this server after you have finalized all of your testing and are ready to initiate a cutover.

When the **Launch cutover instances for X servers** dialog appears, choose **Launch** to begin the cutover.

The AWS Transform MGN console indicates **1 launch job complete** after the cutover has been completed successfully.

This changes your source servers' **Migration lifecycle** status to **Cutover in progress**, indicating that the cutover is in progress but has not yet been finalized.

- **Finalize cutover** – Choose this option to finalize the cutover for this server after you have successfully performed a cutover.

This changes your source servers' **Migration lifecycle** status to **Cutover complete**, indicating that the cutover is complete and that the migration has been performed successfully. In addition, this stops data replication and causes all replicated data to be discarded. All AWS resources used for data replication are terminated.

When the **Finalize cutover for X servers** dialog appears, choose **Finalize**.

The AWS Transform MGN console indicates **X servers cutover. Data replication has been stopped for servers** once the cutover has been completed successfully. The AWS Transform MGN console automatically stops data replication for the cutover source servers to save resource costs. The selected source servers' **Migration lifecycle** column shows the **Cutover** status, the **Data replication** column shows **Disconnected** and the **Next step** column states **Mark as archived**. The source servers have now been successfully migrated into AWS and can be archived.

 Note

This action does not uninstall the AWS Replication Agent from the source server. Use the **Disconnect from service** option under the **Actions** menu when you have completed the migration and want to uninstall the agent from your source server.

- **Revert to "ready for cutover"** – Choose this option to revert a finalized cutover for this server if you encounter any issues or want to reverse the cutover for any reason.

This reverts your source servers' **Migration lifecycle** to the **Ready for cutover** status, indicating that these servers have not undergone cutover.

When the **Revert cutover for X servers** dialog appears, choose **Revert**.

- **Edit launch settings** – Use this option to edit the launch settings for this server. This redirects you to the **Launch settings** tab. [Learn more about Launch settings.](#)

- **Terminate launched instance** – Choose this option if you want to delete your test or cutover instance for any reason at any time. This option can only be selected for a server that has a launched test or cutover instance.

When the **Terminate launched instance** dialog appears, choose **Terminate**.

- **Edit post-launch settings** – Choose this option to edit the post-launch settings for the selected source server or group of source servers. [Learn more about post-launch settings.](#)

Alerts and errors

You can easily distinguish between healthy servers and servers that are experiencing issues on the **Migration dashboard**.

The AWS Transform MGN console is color-coded for ease of use.

- Healthy servers with no errors are characterized by the color blue in both the **Lifecycle** and **Data replication status** boxes.
- Servers that are experiencing temporary issues such as lagging or rescanning, are characterized by the color yellow. These issues do not halt the replication, but may delay it or indicate a bigger problem.
- Servers that are experiencing serious issues such as a loss of connection or a stall are characterized by the color red. You have to fix these issues for data replication to resume. The **Next actions** box provides additional information.

For example, if a stall occurred during initiation, you would scroll down to **Replication initiation steps**, where the problematic step is marked with a red 'x'.

Review source server information

The **Server info** tab displays general server information, hardware, and network information:

- **General information**
 - **Last updated:** when was the data in this tab updated.
 - **Date added:** when was this server added to the service.
 - **AWS ID:** the ID of this source server resource
 - **arn:** the Amazon Resource Name for this source server.

- **Identification hints:** under most circumstances, the source server name is the best identifier, as it is what is used throughout the console as the name of the source server. If you need to validate which external server this is referring to in your data center, you can use one of the additional fields: Fully qualified domain name, VMware virtual machine identifier (only if source is VMware), AWS instance ID (only if source is running on AWS).
- **Hardware and operating system:** the CPUs, RAM, disks, and network interfaces on the external server, as well as the type and full name of the operating system running on that server. The disks shown are all the disks on the source server, and may include disks not being replicated.
- **Recommended instance type:** this is the EC2 instance type the service is auto-recommending to use for the launched recovery instance. This is based only on the CPUs and RAM at the source (and not on utilization information). This is the instance type that is launched for this server by default.

Information shown includes:

- **Last updated**
- **Date added**
- **Hostname**
- **Fully qualified domain name**
- **VMware virtual machine identifier (if relevant)**
- **AWS instance ID**
- **AWS ID**
- **ARN**
- **Operating system** information
- **CPUs**
- **RAM**
- **Primary network interface** Displays the IPv6 address of the interface.
- **Recommended instance type**

Add or remove tags assigned to source servers

The Tags section shows any tags that have been assigned to the server. A tag is a label that you assign to an AWS resource and can be used to search and filter your resources or track your AWS

costs. Each tag consists of a key and an optional value. Learn more about AWS tags in [this Amazon EC2 article](#).

To add tags, take these steps:

- Choose **Manage tags**.
- The **Manage tags** page opens. Choose **Add new tag**.
- Add a tag **Key** and an optional tag **Value**.
- Choose **Save**.

To remove a tag, take these steps:

- Choose **Remove**, located to the right of the tag you want to remove.
- Choose **Save**.

Review disk settings for source servers

The **Disk settings** tab shows a list of all of the disks on the source server and information for each disk.

Disk settings include:

- **Disk name**
- **Staging disk type** – The storage type being used for the disk. When using Amazon EBS, this shows the Amazon EBS volume type. When using FSx for ONTAP, this shows the FSx for ONTAP volume configuration.
- **Replicated storage** – The amount of storage that has been replicated from the disk to the Replication Server.
- **Total storage** – The total storage capacity of the disk.

Change target storage type

You can change the target storage type for all disks on a source server. The available options are Amazon Elastic Block Store (Amazon EBS) and Amazon FSx for NetApp ONTAP. This change applies to all disks on the server.

When using Amazon EBS, you can also customize the Amazon EBS volume type for each individual disk or group of disks. For details, see [Change staging disk type](#) in the [Amazon EBS configuration](#).

For information on configuring FSx for ONTAP as the target storage type, see the [FSx for ONTAP configuration](#).

Edit replication settings for a source server

The **Replication settings** tab allows you to edit the replication settings for an individual source server.

After the source server is added to AWS Transform MGN, the replication settings that are defined in the replication settings template are automatically applied to the server. You can later edit them for a single source server through the **Replication settings** tab.

Edit each setting as required and then choose **Save replication settings**.

[Learn more about replication settings.](#)

Review launch settings for a source server

The launch settings are a set of instructions that comprise an EC2 launch template and other settings, which determine how a test or cutover instance is launched for each source server on AWS.

Launch settings, including the EC2 launch template, are automatically created every time you add a server to AWS Transform MGN.

The launch settings can be modified at any time, including before the source servers have even completed initial sync.

[Learn more about individual launch settings.](#)

Review post-launch settings

Post-launch settings allow you to control and automate actions performed after the server has been launched in AWS. These settings are created automatically based on the **Post-launch settings template**.

You must activate the post-launch actions using one of these options:

- **Activating the post-launch actions for a specific server:**

- Navigate to the **Source servers** page and select a source server.
- Choose **Post-launch settings > Edit**.
- You are redirected to the **Edit post-launch settings** screen. Activate the toggle and choose **Save settings**.

Alternatively, you can select a specific source server, open the **Test and cutover** drop-down menu located in the top right corner of the screen and select **Edit post-launch settings**.

- **Activating the post-launch actions for all servers:**

- Navigate to the **Settings** page and choose **Post-launch settings template**. You only need to do this once and the change applies to all newly added servers.

After the post-launch actions have been activated from the template, you can deactivate and activate them for individual servers. [Learn more about activating post-launch settings](#).

The settings configured in the template are applied to every newly added server. You can change the settings for existing and newly added servers individually within the server details view.

The **Post-launch settings** template allows you to control various post-launch actions, including:

- Deployment of test and cutover instances
- Disaster recovery configuration (installing the AWS Replication Agent for AWS DRS and configuring the target disaster recovery AWS Region)
- Operating system conversion on the target machine
- License and subscription changes on the target machine

Editing the post-launch settings

To edit the post-launch settings for a single source server, check the box to the left of the Hostname of each source server for which you want to edit the post-launch settings, open the **Replication** menu, and choose **Edit post-launch settings**.

Alternatively, when editing the settings for a single server, you can choose **Edit** from the **Post-launch settings** tab.

These settings can be edited within the post-launch settings template. Once you have edited all your settings, choose **Save template**.

Types of post-launch actions

AWS MGN supports post-launch modernization actions, giving you the opportunity to move and improve. All post-launch actions are based on SSM documents (either public or ones you created) that can be executed on your EC2 launch instances.

There are 2 types of post-launch actions:

- **Predefined post-launch actions** – These out-of-the box actions are based on public SSM documents that cannot be changed and have certain unchangeable parameters such as the platform name and order. Fields are prepopulated with the necessary values and only need to be activated or deactivated. For a list of the available actions, see [Predefined post-launch actions reference](#)
- **Custom post-launch actions** – These actions are based on SSM documents that you create and upload to your account. To add a custom post-launch action, see [Create a custom post-launch action](#). To edit a custom post-launch action, see [Edit custom post-launch actions](#).

Use the **Filter by** options on the left-hand side to filter the available actions according to your preferences.

Choose the settings icon in the right-hand corner of the screen to alternate between card and list view, according to your preferences.

Activating and deactivating post-launch actions

This setting controls whether post-launch actions are active or inactive. You must leave the **Install Systems Manager agent and allow executing actions on launched servers** option toggled in order for post-launch actions to work. Untoggling the option disallows AWS MGN to install the SSM Agent on your servers and post-launch actions are no longer executed on them.

The feature is activated and deactivated at the account level from the **Settings > Post-launch template** screen. [Learn more about activating post-launch settings](#).

After it was activated once, the feature can also be deactivated and reactivated for a single server. Simply selecting a server, go to the **Post-launch settings** tab and choose **Edit**.

When the feature is inactive:

- All actions are hidden.

- You are not able to activate actions at the account level or the feature level.

When the feature is active:

- The actions are visible.
- You can activate them.

Deploying post-launch actions

Use this setting to choose whether to deploy the post-launch actions only on your cutover instances or on both cutover and test instances.

Source server actions in AWS Transform MGN

When you install the AWS Replication Agent on a source server in your on-premises environment, AWS Transform MGN begins replicating that server to AWS. Throughout the migration lifecycle, you can perform a variety of actions on your source servers directly from the AWS Transform MGN console. These actions are organized by migration phase: before migration begins, during active replication, and after cutover is complete.

Understanding which actions are available at each phase helps you plan your migration project, minimize downtime, and ensure a successful outcome.

Topics

- [Requirements](#)
- [Actions before migration](#)
- [Actions during migration](#)
- [Actions after migration](#)
- [Source server actions summary](#)

Requirements

Before performing any source server actions, ensure that the following requirements are met:

- AWS Transform MGN has been [initialized](#) in the target AWS Region.
- The AWS Replication Agent has been installed on the source server. [Learn more about adding source servers.](#)
- The IAM user or role performing actions has the **AWSApplicationMigrationFullAccess** managed policy attached, or equivalent permissions.
- Network connectivity between the source server and the MGN service endpoints is maintained throughout the migration lifecycle. [Learn more about network requirements.](#)

Actions before migration

After installing the AWS Replication Agent and before initiating a test or cutover launch, you can perform the following actions on your source servers to prepare them for migration.

Edit replication settings

You can customize the replication settings for each source server individually, or apply settings to multiple servers at once. Replication settings control how data is replicated from the source server to the staging area in AWS, including the replication server instance type, replication subnet, and encryption options.

To edit replication settings for a source server:

1. Open the AWS Transform MGN console and navigate to the **Source servers** page.
2. Select the checkbox next to one or more source servers.
3. Choose **Replication**, then choose **Edit replication settings**.
4. Modify the settings as needed and choose **Save**.

Note

Changes to replication settings take effect on the next replication cycle. Modifying settings while initial sync is in progress may cause the sync to restart.

Edit launch settings

Launch settings control how the migrated server will be launched as an Amazon EC2 instance in AWS. You can configure instance type, subnet, security groups, IAM instance profile, and other launch parameters before initiating a test or cutover launch.

To edit launch settings for a source server:

1. Open the AWS Transform MGN console and navigate to the **Source servers** page.
2. Choose the source server name to open the **Server details** view.
3. Choose the **Launch settings** tab.
4. Modify the settings as needed and choose **Save**.

[Learn more about launch settings.](#)

Add or manage tags

You can add, edit, or remove tags on source servers at any time. Tags help you organize and identify servers, and can be used to filter the source servers list.

To manage tags for a source server:

1. Open the AWS Transform MGN console and navigate to the **Source servers** page.
2. Choose the source server name to open the **Server details** view.
3. Choose the **Tags** tab.
4. Choose **Manage tags** to add, edit, or remove tags. Choose **Save** when done.

Configure post-launch actions

Post-launch actions are scripts or SSM documents that run automatically on the launched instance after a test or cutover launch. You can configure post-launch actions before migration to automate tasks such as software installation, configuration changes, or validation checks.

To configure post-launch actions for a source server:

1. Open the AWS Transform MGN console and navigate to the **Source servers** page.
2. Choose the source server name to open the **Server details** view.
3. Choose the **Post-launch settings** tab.
4. Configure the desired post-launch actions and choose **Save**.

[Learn more about post-launch settings.](#)

Assign server to an application

You can group source servers into applications to manage and migrate related servers together. Assigning servers to an application before migration helps you coordinate test and cutover launches across multiple servers.

To assign a source server to an application:

1. Open the AWS Transform MGN console and navigate to the **Source servers** page.
2. Select the checkbox next to one or more source servers.

3. Choose **Actions**, then choose **Assign to application**.
4. Select an existing application or create a new one. Choose **Assign**.

[Learn more about applications.](#)

Actions during migration

Once the AWS Replication Agent is installed and data replication is active, you can perform the following actions on your source servers. These actions allow you to monitor replication health, manage the replication process, and progress through the migration lifecycle.

Start data replication

Data replication begins automatically after the AWS Replication Agent is installed. However, if replication has been paused or stopped, you can restart it manually.

To start data replication:

1. Open the AWS Transform MGN console and navigate to the **Source servers** page.
2. Select the checkbox next to one or more source servers.
3. Choose **Replication**, then choose **Start data replication**.

Pause data replication

You can pause data replication for a source server. When replication is paused, the replication server in the staging area is stopped, which reduces costs. Data changes on the source server are not replicated while replication is paused.

To pause data replication:

1. Open the AWS Transform MGN console and navigate to the **Source servers** page.
2. Select the checkbox next to one or more source servers.
3. Choose **Replication**, then choose **Pause data replication**.

Important

Pausing replication will cause lag to accumulate on the source server. When replication is resumed, a resync will be required to bring the target up to date. This may extend the time needed before the server is ready for cutover.

Resume data replication

After pausing replication, you can resume it at any time. AWS Transform MGN will perform a resync to replicate any data changes that occurred while replication was paused.

To resume data replication:

1. Open the AWS Transform MGN console and navigate to the **Source servers** page.
2. Select the checkbox next to one or more source servers.
3. Choose **Replication**, then choose **Resume data replication**.

Launch test instance

Once the initial sync is complete and the server reaches the **Ready for testing** state, you can launch a test instance. A test instance is a non-disruptive launch that allows you to validate the migrated server in AWS without affecting the source server or the ongoing replication.

To launch a test instance:

1. Open the AWS Transform MGN console and navigate to the **Source servers** page.
2. Select the checkbox next to one or more source servers in the **Ready for testing** state.
3. Choose **Test and cutover**, then choose **Launch test instances**.
4. Review the launch summary and choose **Launch**.

Note

Launching a test instance does not stop replication. Data continues to replicate from the source server while the test instance is running.

Finalize test and mark as ready for cutover

After validating the test instance, finalize the test to mark the server as ready for cutover. This action terminates the test instance and advances the server's migration lifecycle state.

To finalize a test:

1. Open the AWS Transform MGN console and navigate to the **Source servers** page.
2. Select the checkbox next to one or more source servers.
3. Choose **Test and cutover**, then choose **Finalize test**.
4. Confirm the action by choosing **Finalize test** in the dialog.

Launch cutover instance

When you are ready to perform the final migration, launch a cutover instance. Before launching a cutover instance, stop all services on the source server to ensure that no new data changes occur and that the replication lag is zero.

To launch a cutover instance:

1. Confirm that the data replication lag is zero on the source server.
2. Stop all operational services on the source server.
3. Open the AWS Transform MGN console and navigate to the **Source servers** page.
4. Select the checkbox next to one or more source servers in the **Ready for cutover** state.
5. Choose **Test and cutover**, then choose **Launch cutover instances**.
6. Review the launch summary and choose **Launch**.

Important

Launching a cutover instance is a significant step in the migration process. Ensure that you have completed all acceptance testing and that your team is prepared for the cutover window before proceeding.

Monitor replication health

You can monitor the replication health and progress of each source server from the AWS Transform MGN console. The **Source servers** page displays the current data replication state and migration lifecycle state for each server.

To view detailed replication information for a source server:

1. Open the AWS Transform MGN console and navigate to the **Source servers** page.
2. Choose the source server name to open the **Server details** view.
3. Review the **Migration dashboard** section to see the current lifecycle state, data replication status, and any alerts or errors.

The following data replication states may be displayed:

- **Initial sync** – The agent is performing the first full replication of the source server's data to AWS.
- **Healthy** – Replication is active and up to date with no lag.
- **Lagging** – Replication is active but behind the source server. This is typically a temporary state.
- **Stalled** – Replication has stopped due to an error. Review the alerts and take corrective action.
- **Paused** – Replication has been manually paused.
- **Disconnected** – The agent on the source server has lost connectivity to AWS Transform MGN.

Actions after migration

After a successful cutover, you can perform the following actions to complete the migration and clean up resources.

Finalize cutover

After confirming that the cutover instance is running correctly in AWS, finalize the cutover. This action marks the migration as complete for the source server and stops data replication.

To finalize a cutover:

1. Validate that the cutover instance is functioning correctly in AWS by testing connectivity and application behavior.
2. Open the AWS Transform MGN console and navigate to the **Source servers** page.

3. Select the checkbox next to one or more source servers.
4. Choose **Test and cutover**, then choose **Finalize cutover**.
5. Confirm the action by choosing **Finalize cutover** in the dialog.

Important

Finalizing a cutover stops data replication and terminates the replication server in the staging area. This action cannot be undone. Do not finalize the cutover until you have confirmed that the cutover instance is working as expected.

Disconnect from AWS Transform MGN

After finalizing the cutover, you can disconnect the source server from AWS Transform MGN. This action removes the server from active management in AWS Transform MGN and stops any remaining replication activity. The source server itself is not affected.

To disconnect a source server:

1. Open the AWS Transform MGN console and navigate to the **Source servers** page.
2. Select the checkbox next to one or more source servers.
3. Choose **Actions**, then choose **Disconnect from service**.
4. Confirm the action in the dialog.

Note

Do not disconnect the source server from AWS Transform MGN until the cutover instance in AWS is confirmed to be working correctly. You may need to revert to the source server if issues are discovered after cutover.

Archive source server

After the migration is complete and the cutover instance is confirmed to be working correctly, you can archive the source server. Archiving removes the server from the active **Source servers** list and moves it to the archived view. This helps keep your console organized after large-scale migrations.

To archive a source server:

1. Open the AWS Transform MGN console and navigate to the **Source servers** page.
2. Select the checkbox next to one or more source servers.
3. Choose **Actions**, then choose **Archive server**.
4. Confirm the action in the dialog.

 Note

Archived servers can be viewed by toggling the **Show archived servers** option in the **Preferences** menu on the Source servers page.

Delete source server

You can permanently delete a source server from AWS Transform MGN. This removes all data associated with the server from the AWS Transform MGN console, including replication history and settings. This action does not affect the source server itself or any launched Amazon EC2 instances.

To delete a source server:

1. Open the AWS Transform MGN console and navigate to the **Source servers** page.
2. Select the checkbox next to one or more source servers.
3. Choose **Actions**, then choose **Delete server**.
4. Confirm the deletion by entering the server name in the confirmation dialog and choosing **Delete**.

 Important

Deleting a source server is a permanent action and cannot be undone. Ensure that the migration is complete and the cutover instance is working correctly before deleting the source server record from AWS Transform MGN.

Source server actions summary

The following table summarizes the actions available on source servers at each phase of the migration lifecycle.

Source server actions by migration phase

Action	Migration phase	Description
Edit replication settings	Before migration	Configure how data is replicated to the AWS staging area.
Edit launch settings	Before migration	Configure how the migrated server will be launched as an Amazon EC2 instance.
Add or manage tags	Before migration	Organize and identify source servers using tags.
Configure post-launch actions	Before migration	Set up scripts or SSM documents to run after a test or cutover launch.
Assign to application	Before migration	Group related servers into an application for coordinated migration.
Start data replication	During migration	Start or restart data replication from the source server to AWS.
Pause data replication	During migration	Temporarily stop replication to reduce costs. Lag will accumulate.
Resume data replication	During migration	Resume replication after a pause. A resync will be performed.
Launch test instance	During migration	Launch a non-disruptive test instance to validate the migrated server.
Finalize test	During migration	Mark the test as complete and advance the server to the ready-for-cutover state.

Action	Migration phase	Description
Launch cutover instance	During migration	Perform the final migration by launching the cutover instance in AWS.
Monitor replication health	During migration	View replication status, lag, and alerts for each source server.
Finalize cutover	After migration	Mark the migration as complete and stop data replication.
Disconnect from service	After migration	Remove the server from active AWS Transform MGN management.
Archive server	After migration	Move the server to the archived view to keep the console organized.
Delete server	After migration	Permanently remove the server record from AWS Transform MGN.

Associate a Group of Servers with an Application

Many customers have clusters of servers with dependencies between them. AWS Transform MGN provides the user with a way to represent a group of servers by associating them with an **Application**.

You can monitor the migration status and progress of an application and its associated servers. You can also perform operation on the application, such as edit, tagging, archive, as well as bulk operations on the servers associated with the application.

Topics

- [Manage applications](#)
- [Access details for an application](#)

Manage applications

The **Applications** page lists all the applications that have been added to AWS Transform MGN. The **Applications** page allows you to manage your applications and perform a variety of commands for one or more applications (such as controlling replication and launching test and cutover instances).

Interacting with the Applications page

The **Applications** page shows a list of applications. Each row on the list represents a single application.

The **Applications** page provides key information for each application under each of the columns on the page.

The columns include:

- **Selector column** – This blank checkbox selector column allows you to select one or more applications. When an application is selected, you can interact with the application through the **Actions** menu, **Edit**, and **Delete** buttons. Selected applications are highlighted.
- **Application name** – This column shows the unique application name for each application.
- **Wave name** – This column shows the name of the wave the application is associated with. An application cannot be associated with more than one wave at a time.

This column is hidden by default.

- **Migration status** – This column shows the migration status for each application.
 - **Not started** – None of the application associated servers has started replication yet.
 - **In progress** – At least one of the application associated servers has started replication and not all of its servers completed migration.
 - **Completed** – All the application associated servers completed migration (have been cut over).
- **Alerts** – This column shows whether any alerts exist for the application.
 - **Stalled** – An application that has at least one server that is experiencing significant issues, such as a stall.
 - **Lagging** – An application that has at least one server that is experiencing a temporary issue such as lag or backlog.
 - **Healthy** – A healthy active application.

Archived applications do not display any alerts.

- **Number of servers** – This column shows the total number of servers associated with each application.

Topics

- [Add application](#)
- [Edit application](#)
- [Delete application](#)
- [Manage applications](#)
- [Filtering the Applications page](#)

Add application

To add an application, choose **Add application**. When the **Add application** prompt opens, configure the application name, add a description (optional), associate source servers (optional), and add tags (optional).

- **Application name** – Application name is mandatory, with a limit of 256 characters. The name must be unique per account per region. Uniqueness verification for application name in Migration Application Service is case-insensitive.

- **Description** – Application description is optional, with a limit of 600 characters.
- **Servers** – You can add up to 200 servers to an application. Checking a server in the drop-down list will associate it with the application.
- **Tags** – You can add up to 50 tags to an application.

When you are done configuring your application settings, choose **Add application** to create the application.

Edit application

To edit an application, choose **Edit**. When the **Edit application** prompt opens, edit the application name, description, and tags, as well as associate or disassociate source servers.

- **Application name** – Application name is mandatory, with a limit of up to 256 characters. The name must be unique per account per region. Uniqueness verification for application name in Migration Application Service is case-insensitive.
- **Description** – Application description is optional, with a limit of 600 characters.
- **Servers** – You can add up to 200 servers to an application. Checking a server in the dropdown list will associate it with the application. Unchecking an associated server will disassociate it from the application.
- **Tags** – You can add up to 50 tags to an application.

To finalize your changes, choose **Save changes**.

Delete application

To delete an application, choose **Delete**. When the **Delete application** prompt opens, verify that you want to delete the selected application.

Deleting the application will disassociate the servers from the application, but will not delete them.

Choose **Delete** to confirm the deletion.

Manage applications

The **Actions** menu allows you to perform actions on selected applications.

 Note

An application must have **all** of its associated servers in the correct lifecycle for the desired action, otherwise it will be excluded.

The **Actions** menu allows you to perform the following actions:

- **Launch test instances** – Choose this option to launch test instances for this application servers.
- **Mark as "Ready for cutover"** – Choose this option to finalize testing for this application after you have completed all the necessary tests in preparation for cutover.

The **Mark servers as "Ready for cutover"** dialog will appear. Select whether you want to terminate the launched instances used for testing. It is recommended to terminate these instances, as you will be charged for them even though you will no longer need them. Check the **Yes, terminate launched instances (recommended)** box and choose **Continue**.

- **Revert to "ready for testing"** – Choose this option to revert a finalized test for this application if you want to run further tests before initiating a cutover.

The **Revert testing** dialog will appear. Select whether you want to terminate the launched instances used for testing. It is recommended to terminate these instances, as you will be charged for them even though you will no longer need them. Check the **Yes, terminate launched instances (recommended)** box and choose **Revert**.

- **Launch cutover instances** – Choose this option to launch cutover instances for this application servers after you have finalized all of your testing and are ready to initiate a cutover.
- **Finalize cutover** – Choose this option to finalize the cutover for this application servers after you have successfully performed a cutover.

 Note

This action does not uninstall the AWS Replication Agent from the source servers. When you have completed the migration and want to uninstall the agent from your source servers, go to **Source servers** page and select the relevant servers. Use the **Disconnect from service** option under the **Actions** menu.

- **Revert to "ready for cutover"** – Choose this option to revert a finalized cutover for this application if you encounter any issues or want to reverse the cutover for any reason.

- **Start data replication** – Choose this option to start replicating the application source servers.

 Note

This action is applicable if all the application associated servers are **Agentless snapshot based** and are in **Discovered** lifecycle state.

- **Add applications to wave** – Choose this option to associate the selected applications to a wave.
- **Archive applications** – Choose this option to archive the selected applications. You should only archive applications for which you have already performed a cutover.

 Important

An application can be archived only if all servers that compose it are in one of these states: archived, cutover, or disconnected. If that is the case, the application will be archived and the servers that are not yet archived (but can be) will also be archived.

Archived applications will be removed from the main applications page, but can still be accessed through the selector options.

Filtering the Applications page

You can customize the **Applications** page through filtering. Filtering allows you to easily filter your applications by one or multiple properties.

Choose within the **Filter applications** field and choose the filtering property from the **Properties** menu.

You can filter by a variety of properties, including:

- Application name – Filter by application name.
- Application ID – Filter by application ID.
- Wave name – Filter by wave name.
- Migration status – Filter by the migration status (Not started, In progress, Completed).
- Alerts – Filter by health status alert (Stalled, Lagging, Healthy).
- Number of servers – Filter by a number of servers.

- **Tags** - Filter by tags. Relevant specific tag values will appear under the **Tags** category. Choose the value by which to filter.


You can filter by multiple properties at once to narrow down your results.

To clear the selected filtering properties, choose **Clear filters**.

Access details for an application


There are several ways you can access the **Application details** view.

Choose the **Application name** of any application on the **Applications** page.



☐	Application name	Wave name	Migration status	Alerts	Number of servers
☐	App 3	wave 3	In progress	Healthy	2

Choose the **Application** of any server on the **Source servers** page.



☐	Source server name	Application	Alerts	Migration lifecycle	Data replication status	Last snapshot	Next step
☐	ip-172-31-35-46.eu-central-1.compute.internal	arch-app	-	Ready for cutover	Healthy	an hour ago	Launch cutover instance

Choose the **Application name** in the **Server info** tab.

[Migration dashboard](#) | [Server info](#) | [Tags](#) | [Disks settings](#) | [Replication settings](#) | [Launch settings](#) | [Post-launch settings](#)

Server info [Info](#)

General information

Last updated
September 29, 2022 at 10:05 (UTC+3:00)

Date added
September 01, 2022 at 15:40 (UTC+3:00)

AWS ID
s-0d1b696ed7bee64ac

ARN
arn:aws:mgn:eu-west-1:027888403785:source-server/s-0d1b696ed7bee64ac

Application name
[App 3](#)

Hardware

CPU
Model: Intel(R) Xeon(R) CPU E5-2676 v3 @ 2.40GHz, Cores: 1

Identification hints

Hostname
ip-172-31-36-230.eu-central-1.compute.internal

Fully qualified domain name
ip-172-31-36-230.eu-central-1.compute.internal

VMware virtual machine identifier
-

AWS instance ID
i-03df59d8c6e5a59e9

Primary network interface
IP: 172.31.36.230
MAC: 06-A9-1B-34-14-A4

Operating system

Type
Linux

Choose the **Application name** of any application in the **Applications** table inside **Wave details** -> **Applications** tab.

[Applications](#) | [Source servers](#) | [Tags](#)

Alerts [Info](#)

Filter applications



Healthy
1 application, 100%

Migration status [Info](#)

Filter applications




In progress
1 application, 100%

Applications (1) [Info](#)

Application name	Migration status	Alerts	Number of servers	Archived
App 3	In progress	Healthy	2	Active

Last update: September 29, 2022 at 11:08 (UTC+3:00)

The **Application details** view shows information and options for an individual application. Here, you can control and monitor the individual application.



You can also perform a variety of actions on the application, and perform batch operations such as launch Test and Cutover instances for the servers associated with the application.

The **Application details** view is divided into several dashboards:

Topics

- [Review overall application status](#)
- [Source server migration metrics](#)
- [Review application source servers](#)
- [Review application tags](#)

Review overall application status

The **Overview** dashboard provides an overview of the overall application status, including:

- **Description** – The description of the application.
- **State** – The state of the application. **State** can be in one of two states: **Active** or **Archived**.

- **Last status update** – Time stamp of when application status was updated (update occurs every five minutes).
- **Wave name** – Name of the wave that the application is associated with.
- **Migration status** – The application migration status.

Application **Migration status** can have one of the following values:

Not started – If none of its servers has started replication yet.

Completed – If all of its servers completed migration (have been cutover).

In progress – At least one of its servers has started replication and not all of its servers completed migration.

- **Alerts** – The application alert.

An application that has at least one server that is experiencing significant issues, such as a stall, will display a **Stalled** status.

An application that has at least one server that is experiencing a temporary issue such as lag or backlog will display a **Lagging** status.

A healthy active application will display a **Healthy** status.


An archived application will not display a status.

Source server migration metrics

The source server migration metrics show an aggregated overview of the application associated servers on three topics: **Alerts**, **Data replication status** and **Migration status**.

Review source server alerts


The source server **Alerts** migration metric presents an aggregated overview of the application associated servers alerts. You can look up an individual source server **Alerts** status at the **Source servers** table at the bottom of the page.



- A healthy server for which a test or cutover instance has not been launched will display a **Healthy** status.
- A healthy server for which a test or cutover instance has been launched will display a **Healthy** status.
- A server that is experiencing a temporary issue such as lag or backlog will display a **Lagging** status.
- A server that is experiencing significant issues, such as a stall, will display a **Stalled** status.

Review source server data replication status

The source server **Data replication status** migration metric presents an aggregated overview of the application associated servers data replication status. You can look up an individual source server's **Data replication status** at the **Source servers** table at the bottom of the page.




Source server **Data replication status** can have one of the following values:

- **Transferring snapshot**
- **Initial sync**
- **Finalizing sync**
- **Lagging**
- **Healthy**
- **Stalled**
- **Rescanning**
- **Not started**
- **Initiating**
- **Creating snapshot**
- **Paused**
- **Disconnected**

Review the application source server migration lifecycle

The source server **Migration lifecycle** metric shows an aggregated overview of the application associated servers migration lifecycle. You can look up an individual source server **Migration lifecycle** status at the **Source servers** table at the bottom of the page.



Source server **Migration lifecycle** can have one of the following values:

- **Stopped**
- **Not ready**
- **Ready for testing**
- **Test in progress**
- **Ready for cutover**
- **Cutover in progress**
- **Cutover complete**
- **Disconnected**
- **Discovered**

Review application source servers

The **Source servers** table lists all the servers that are associated with the application.

To perform batch operations on all the servers, use the application **Actions** menu at the top of the page. To perform an operation on a single server, go to the specific server's **Server details** page by choosing the server **Source server name**.

Review application tags

The **Tags** section shows any tags that have been assigned to the application. A tag is a label that you assign to an AWS resource. Each tag consists of a key and an optional value. You can use tags to search and filter your resources or track your AWS costs. Learn more about AWS tags in [this Amazon EC2 article](#).

Group Source Servers and Applications in Waves

AWS Transform MGN helps users manage their migration by grouping **Source servers** and **Applications** in **Waves**. These are logical groups, describing the migration plan over time.

You can:

- Monitor the wave's migration status, progress, and associated applications
- Perform operations on the wave, such as editing, tagging, and archiving
- Perform bulk operations on the applications associated with the wave

Topics

- [Manage migration waves](#)
- [Review migration wave details](#)

Manage migration waves

The **Waves** page lists all the waves that have been added to AWS Transform MGN. The **Waves** page allows you to manage your waves and perform a variety of commands for one or more waves (such as controlling replication and launching test and cutover instances).

Interacting with the Waves page

The **Waves** page shows a list of waves. Each row on the list represents a single wave.

The **Waves** page provides key information for each wave under each of the columns on the page.

The columns include:

- **Selector column** – This blank checkbox selector column allows you to select one or more waves. When a wave is selected, you can interact with the wave through the **Actions** menu, **Edit**, and **Delete** buttons. Selected waves are highlighted.
- **Wave name** – This column shows the unique wave name for each wave.
- **Migration status** – This column shows the migration status for each wave.
 - **Not started** – If none of the wave's associated servers has started replication yet.

- **In progress** – At least one of the wave's associated applications has started replication and not all of its applications completed migration.
- **Completed** – If all the wave associated applications completed migration (have been cutover).
- **Alerts** – This column shows whether any alerts exist for the wave.

A wave that has at least one application that is experiencing significant issues, such as a stall, will display a **Stalled** status.

A wave that has at least one application that is experiencing a temporary issue such as lag or backlog will display a **Lagging** status.

A healthy active wave will display a **Healthy** status.

Archived waves do not display any alerts.

- **Number of applications** – This column shows the total number of applications associated with each wave.

Topics

- [Add wave](#)
- [Edit wave](#)
- [Delete wave](#)
- [Manage selected waves](#)
- [Filtering migration waves](#)

Add wave

To add a wave, choose the **Add wave** button. You will then see the **Add wave** prompt that includes the following parameters:

- **Wave name** – Select a wave name. The name must be unique per account per region. Note that uniqueness verification for wave name in Migration Application Service is case-insensitive.
- **Description (optional)** – Add a description of the wave according to your preferences.
- **Associate applications (optional)** – You can add up to 200 applications to a single wave. Checking an application in the dropdown list will associate it with the wave.
- **Add tags (optional)** – You can add up to 50 tags according to your preferences.

Choose **Add wave** to create the wave.

Edit wave

To edit a wave, choose **Edit wave**. You will see the **Edit wave** prompt, which allows you to edit the following parameters:

- **Wave name** – Select a wave name. The name must be unique per account per region. Note that uniqueness verification for wave name in Migration Application Service is case-insensitive.
- **Description (optional)** – Add a description of the wave according to your preferences.
- **Associate applications (optional)** – You can add up to 200 applications to a single wave. Checking an application in the dropdown list will associate it with the wave.
- **Add tags (optional)** – You can add up to 50 tags according to your preferences.

After you edit the parameters as required, choose the **Save changes** button to save your changes.

Delete wave

To delete a wave, choose the **Delete wave** button and the **Delete wave** prompt will open. You will need to verify that you want to delete the selected wave.

When you delete the wave, all associated applications will disassociate from the wave but will not be deleted.

Choose **Delete** to confirm the deletion.

Manage selected waves

The **Actions** menu allows you to perform actions on selected waves.

Note

A wave must have **all** of its associated servers in the correct lifecycle for the desired action, otherwise it will be excluded.

Use this menu to perform the following actions:

- **Launch test instances** – Choose this option to launch test instances for this wave servers.
- **Mark as "Ready for cutover"** – Choose this option to finalize testing for this wave after you have completed all the necessary tests in preparation for cutover.

After you choose this option, the **Mark servers as "Ready for cutover"** dialog appears. Select whether you want to terminate the launched instances used for testing. It is recommended to terminate these instances, as you will be charged for them even though you will no longer need them. Check the **Yes, terminate launched instances (recommended)** box and choose **Continue**.

- **Revert to "ready for testing"** – Choose this option to revert a finalized test for this wave if you want to run additional tests before initiating a cutover.

The **Revert testing** dialog will appear. Select whether you want to terminate the launched instances used for testing. It is recommended to terminate these instances, as you will be charged for them even though you will no longer need them. Check the **Yes, terminate launched instances (recommended)** box and choose **Revert**.

- **Launch cutover instances** – Choose this option to launch cutover instances for this wave servers after you have finalized all of your testing and are ready to initiate a cutover.
- **Finalize cutover** – Choose this option to finalize the cutover for this wave servers after you have successfully performed a cutover.

 Note

This action does not uninstall the AWS Replication Agent from the source servers. When you have completed the migration and want to uninstall the agent from your source servers, go to **Source servers** page and select the relevant servers. Use the **Disconnect from service** option under the **Actions** menu.

- **Revert to "ready for cutover"** – Choose this option to revert a finalized cutover for this wave if you encounter any issues or want to reverse the cutover for any reason.
- **Start data replication** – Choose this option to start replication of the wave source servers.

 Note

This action is applicable if all the wave's associated servers are **Agentless snapshot based** and are in **Discovered** lifecycle state.

- **Archive waves** – Choose this option to archive the selected waves. You should only archive waves for which you have already performed a cutover.

Important

A wave can be archived only if all servers that are part of it are in one of these states: archived, cutover, or disconnected. If that is the case, the wave and its associated applications will be archived. The servers that are not yet archived (but can be) will also be archived.

- Archived waves will be removed from the main Waves page, but can still be accessed through the selector options.

Filtering migration waves

Use filtering to easily filter your waves by one or multiple properties.

Choose within the **Filter waves** field and choose the filtering property from the **Properties** menu.

You can filter by a variety of properties, including:

- Wave name – Filter by wave name.
- Wave ID – Filter by wave ID.
- Migration status – Filter by the migration status (Not started, In progress, Completed).
- Alerts – Filter by health status alert (Stalled, Lagging, Healthy).
- Number of applications – Filter by a number of applications.
- Tags – Filter by tags. Relevant specific tag values will appear under the **Tags** category. Choose the value to filter by.

You can filter by multiple properties at once to narrow down your results.

To clear the selected filtering properties, choose **Clear filters**.

Review migration wave details

There are several ways you can access the **Wave details** view.

Choose the **Wave name** of any wave on the **Waves** page.

Waves (2) [Info](#)

Actions ▾

Edit

Delete

Add wave

Active waves ▾

🔍

Filter waves by property or value

<

1

>

🔗

☐	Wave name	▲	Migration status	▾	Alerts	▾	Number of applications	▾
☐	wave 1	🕒	In progress		🟢	Healthy	1	

Choose the **Wave** of any application on the **Applications** page.

Applications (3) Info						Actions Edit Delete Add application	
Active applications		Filter applications by property or value			< 1 > ⚙		
☐	Application name	Wave name	Migration status	Alerts	Number of servers		
☐	App 1	Wave 2	In progress	Healthy	3		

Choose the **Wave name** in the **Overview** dashboard inside **Application details** of an application.

App 2 (app-083181d1ef5a1ea8d)			Actions Edit Delete	
Overview Info				
Description	Last status update		Migration status	
App 2 description	October 11, 2022 at 13:06 (UTC+3:00)		In progress	
State	Wave name		Alerts	
Active	wave 1		Healthy	

The **Wave details** view shows information and options for an individual wave. Here, you can control and monitor the individual wave.

Application Migration Service > Waves > MGN-6583

MGN-6583 (wave-322545120585b477f) Actions ▾ Edit Delete

Overview [Info](#)

Description This is the wave description.	Wave start time September 29, 2022 at 12:02 (UTC+3:00)	Migration status ☹ In progress
State Active	Current duration 13 days	Alerts ✔ Healthy
Last status update October 12, 2022 at 09:53 (UTC+3:00)		

[Applications](#) | [Source servers](#) | [Tags](#)

Alerts [Info](#)

Filter applications

Select status to filter ▾

Status	Count	Percentage
Healthy	1	33.3%
Archived	2	66.7%

Migration status [Info](#)

Filter applications

Select status to filter ▾

Status	Count	Percentage
In progress	1	33.3%
Completed	2	66.7%

You can also perform a variety of actions on the wave, and perform batch operations such as launch test and cutover instances for the servers associated with the wave.

The **Wave details** view is divided into several dashboards:

Topics

- [Review overall wave status](#)
- [Applications](#)
- [Source servers](#)

Review overall wave status

The **Overview** dashboard provides an overview of the overall wave status, including:

- **Description** – The description of the wave.
- **State** – The state of the wave. **State** can be in one of two states: **Active** or **Archived**

- **Last status update** – Time stamp of when wave status was updated (update occurs every five minutes).
- **Wave start time** – Time stamp of when the earliest replication started for a server associated with this wave.
- **Current duration** – Duration of replication time since **Wave start time**. If wave is archived, duration is until the moment the wave was archived.
- **Migration status** – The wave migration status.

Wave **Migration status** can have one of the following values:

- **Not started** – If none of its applications has started replication yet.
- **Completed** – If all of its applications completed migration (have been cutover).
- **In progress** – At least one of its applications has started replication and not all of its applications completed migration.
- **Alerts** – The wave alert.
 - A wave that has at least one application that is experiencing significant issues, such as a stall, will display a **Stalled** status.
 - A wave that has at least one application that is experiencing a temporary issue such as lag or backlog will display a **Lagging** status.
 - A healthy active wave will display a **Healthy** status.
 - An archived wave will not display a status.

Applications

The **Applications** tab shows migration metrics aggregating statuses as well as a list of all the wave associated applications.


The application migration metrics show an aggregated overview of the wave's associated servers on two topics: **Alerts** and **Migration status**.

Topics

- [Review alerts on applications in a wave](#)
- [Review the migration status of applications in a wave](#)
- [Review all applications associated with a wave](#)

Review alerts on applications in a wave


The application **Alerts** metric provides an aggregated overview of the alerts related to the wave's associated applications. You can look up an individual application **Alerts** status in the **Applications** table.



- A healthy application will display a **Healthy** status.
- An application that is experiencing a temporary issue such as lag or backlog will display a **Lagging** status.
- An application that is experiencing significant issues, such as a stall, will display a **Stalled** status.

Review the migration status of applications in a wave

The application **Migration status** metric provides an aggregated overview of the migration status of the wave's associated applications. You can look up an individual application's **Migration status** in the **Applications** table.



Application **Migration status** can have one of the following values:

- **Not started**
- **In progress**
- **Completed**

Review all applications associated with a wave

The **Applications** table lists all the applications that are associated with the wave.

You can perform batch operations on all the applications via the wave **Actions** menu at the top of the page. You can perform an operation on a single application from its own **Application details** page, by choosing the application's **Application name**.

Source servers

The **Source servers** tab shows migration metrics aggregating statuses as well as a list of all the wave's associated source servers.


The source server migration metrics provide an aggregated overview of the wave's associated servers on three topics: **Alerts**, **Data replication status**, and **Migration status**.

Topics

- [Review alerts on source servers in a wave](#)
- [Review data replication status of the servers in a wave](#)
- [Review the migration lifecycle for the servers in a wave](#)
- [Review the source servers in a wave](#)
- [Review tags assigned to a wave](#)

Review alerts on source servers in a wave


The source server **Alerts** metric provides an aggregated overview of the alerts related to the wave's associated servers. You can look up an individual source server **Alerts** status in the **Source servers** table.



- A healthy server for which a test or cutover instance has not been launched will display a **Healthy** status.
- A healthy server for which a test or cutover instance has been launched will display a **Healthy** status.
- A server that is experiencing a temporary issue such as a lag or backlog will display a **Lagging** status.
- A server that is experiencing significant issues, such as a stall, will display a **Stalled** status.

Review data replication status of the servers in a wave

The source server **Data replication status** metric provides an aggregated overview of the data replication status of the wave's associated servers. You can look up an individual source server's **Data replication status** in the **Source servers** table.




Source server's **Data replication status** can have one of the following values:

- **Transferring snapshot**

- **Initial sync**
- **Finalizing sync**
- **Lagging**
- **Healthy**
- **Stalled**
- **Rescanning**
- **Not started**
- **Initiating**
- **Creating snapshot**
- **Paused**
- **Disconnected**

Review the migration lifecycle for the servers in a wave

The source server's **Migration lifecycle** metric provides an aggregated overview of the migration lifecycle of the wave's associated servers . You can look up an individual source server's **Migration lifecycle** status in the **Source servers** table.



The source server's **Migration lifecycle** can have one of the following values:

- **Stopped**
- **Not ready**
- **Ready for testing**
- **Test in progress**
- **Ready for cutover**
- **Cutover in progress**
- **Cutover complete**
- **Disconnected**
- **Discovered**

Review the source servers in a wave

The **Source servers** table lists all the servers that are associated with the wave.

You can perform batch operations on all the servers via the wave **Actions** menu at the top of the page. You can perform an operation on a single server from its own **Server details** page, by choosing the server **Source server name**.

Review tags assigned to a wave

The **Tags** section shows any tags that have been assigned to the wave. A tag is a label that you assign to an AWS resource. Each tag consists of a key and an optional value. You can use tags to search and filter your resources or track your AWS costs. Learn more about AWS tags in [this Amazon EC2 article](#).

MGN Connectors

MGN connectors enable you to run commands on multiple source servers. Use connectors for large-scale migrations with multiple operating system types and versions, which may be distributed across multiple data centers. A connector can help:

- Verify the prerequisites are met for installation of the MGN replication agent on the source servers.
- Install the MGN replication agents on the source servers.

This feature, combined with the post-launch action framework, offers automation across the entire deployment process.

Topics

- [Prerequisites for installing the MGN connector](#)
- [Architecture overview for MGN connector](#)
- [Set up the MGN Connector](#)
- [IAM roles needed for the MGN connector](#)
- [Installing the MGN connector on a secured network](#)
- [Review details about your MGN connectors](#)
- [Manage your MGN Connectors](#)

Prerequisites for installing the MGN connector

To use the MGN connector you must meet these prerequisites.

General prerequisites

- While the MGN connector can be deployed on the same server that hosts the MGN vCenter Client installer (agentless appliance), we recommend that the MGN connector be installed on a dedicated server.
- The *openssl* library must be installed on the server.
- You must have the required [permissions](#).

Note

The MGN connector is not supported for IPv6.

Operating systems that support the MGN connector

The MGN connector can be installed on servers running the following Linux versions:

- Ubuntu 18.x+ (64 bit) - 22.04 (x86_64)
- Amazon Linux 2 (x86_64)
- RHEL8.x (x86_64)

SSM agent installation requirements

Installation of the MGN Connector also installs the SSM agent.

- If the SSM agent is already installed on the server you must uninstall it before installing the MGN connector. See [Uninstalling SSM Agent from Linux instances](#) in the *AWS Systems Manager User Guide*.
- A minimum of 200 MB of free disk space and 200 KB of free disk space in the `/var` directory.


Security recommendations for MGN connector

We recommend that the MGN connector server is only accessed by authorized personnel and has all the required OS patches. We also recommend that the servers to which the MGN connector connects have all the required OS patches.

If you configure [outputting logs to S3](#), you will first [create an Amazon S3 bucket](#). We recommend that you apply S3 bucket [S3 security practices](#)

Architecture overview for MGN connector

The following diagram describes the architecture overview when using AWS Transform MGN with MGN connector.



Set up the MGN Connector

To add an MGN connector, click **Add MGN connector**, to open the Add MGN connector page. Set up your MGN connector by providing the following:

1. **Connector name:** The MGN connector name is used to identify the connector. This field is mandatory, and limited to 256 characters. The name must be unique (case-insensitive) per account per Region.
2. **IAM roles** - The MGN console will automatically create the required IAM roles in your account. For details about the roles created, see [Create roles using the MGN console](#).
 - **Automatic role generation:** The setup page will automatically generate these two roles in your account:
 - **AWSApplicationMigrationConnectorManagementRole** - Used during agent installation to access credentials.

- **AWSApplicationMigrationConnectorSharingRole_<ACCOUNT-ID>** - Contains permissions for agent installation.
 - **IAM roles deployment scope:**
 - **Individual account:** For an MGN connector in an individual account, the roles are created automatically.
 - **Multiple accounts:** If the MGN connector manages source servers from multiple accounts, set up the global view feature and set up your AWS Organization, as described in [Manage large-scale migrations with global view](#). Alternatively, you can download a CloudFormation template from the setup page to deploy the IAM roles yourself.
3. **SSM Hybrid Activation** - Create a 30-day SSM Hybrid Activation for secure communication between the connector and AWS Systems Manager. This activation enables the connector to register as a managed instance.
 4. **Acknowledge and create resources** - Click to generate the selected resources.
 5. **Installation command** - The setup page generates a one-line installation command with all necessary credentials and configuration.
 6. **Connector installation** - Install the connector on a Linux machine in your environment. For information on the required Linux machine, see [Prerequisites](#).
 - a. Copy the installation link from the setup page.
 - b. SSH into your chosen Linux machine.
 - c. Paste and execute the installation command.
 - d. Wait for installation to complete (typically 2–3 minutes).
 7. **Register servers** - Go to the MGN Connector page in the AWS Transform MGN console and click on the MGN connector name. Go to "Register servers" to attach source servers with the MGN connector.
 8. **Configure credentials** - Select the source servers on which you want to install the replication agent. Go to **Actions** and then **Register server credentials**. You can use an existing AWS Secrets Manager secret or create a new one. If multiple source servers share the same secret, you can select them together and apply the secret to all of them in a single operation. For more information, see [Register server credentials](#).
 9. **Agent deployment** - Once credentials are configured and verified, select the source servers and go to **Actions**, then choose **Install replication agent**. You can keep track of the agent installation status in the **Agent installed** column in the console, or on the **Command history** page.

The deployment process for each server:

- a. MGN sends deployment commands to the connector via SSM.
- b. The connector retrieves credentials from AWS Secrets Manager.
- c. The connector connects to the source server using the configured credentials.
- d. The connector validates that the source server meets all prerequisites required to run the replication agent.
- e. The connector installs and configures the replication agent.
- f. The connector verifies successful installation and connectivity.

Connector reuse and lifecycle

When deploying agents for subsequent waves, you can use an existing connector or create a new one. MGN connector console page lists all connectors configured in your account, showing the connector name, attached server count, and last seen date.

SSM Hybrid Activations expire after 30 days. The activation is required only for installing the connector on the Linux machine. Once the connector is installed, you can continue to use it to install replication agents on source servers even after the activation expires. If you need to install the connector on a new machine after the activation has expired, you need to create a new connector through the setup process.

IAM roles needed for the MGN connector

To use MGN connector you must have these required IAM roles for individual accounts and AWS Organizations networks:

- **AWSApplicationMigrationConnectorManagementRole**
- **AWSApplicationMigrationConnectorSharingRole_management-account-id** - Needed in *every* account that contains source servers on which the connector installs agents. MGN supports cross-account scenarios. In an individual account setup this is the connector's own account. In an organization, the role is needed on *every* account, including the management account. *management-account-id* is the ID of the account in which the connector is created.

Note

The **MGNConnectorInstallerRole** is no longer required. Its permissions are now included in the **AWSApplicationMigrationConnectorManagementRole**, whose credentials the connector installer obtains from the AWS Systems Manager agent through the SSM hybrid activation. If you created the **MGNConnectorInstallerRole** previously, it is no longer used.

You can create these roles in the following ways:

- **Using the AWS Transform MGN console (recommended):** When you add a connector using the MGN console, the service can create the required roles for you, in an individual account or across all member accounts of your AWS Organization. For details about the roles the console creates, see [Create roles using the MGN console](#).
- **Manually, for an individual account:** For an MGN connector in an individual account, create these roles as described in [Create roles manually](#).
- **With a CloudFormation StackSet, for multiple accounts:** If the MGN connector manages source servers from multiple accounts, set up the global view feature and set up your AWS Organization, as described in [Manage large-scale migrations with global view](#). After you set up your AWS Organization, configure the CloudFormation StackSet to create the **AWSApplicationMigrationConnectorSharingRole_management-account-id** role in member accounts. For instructions, see [Deploy role using a CloudFormation template](#).

Create roles using the MGN console

When you add an MGN connector using the MGN console, MGN can create the required IAM roles for you. On the **MGN connectors** page, choose **Add connector**, and then select **Create IAM Roles**. MGN creates the following roles:

- **AWSApplicationMigrationConnectorManagementRole** - created in the current account. This is the role that is assumed by the connector.
- **AWSApplicationMigrationConnectorSharingRole_management-account-id** - created in the current account, where *management-account-id* is the ID of the account in which the connector is created. This role is assumed by the connector and by the MGN service to perform actions on source servers, for example, to install the AWS Replication Agent. MGN supports cross-account scenarios: this role must exist in *every* account that contains source servers on which

the connector installs agents. If you choose to deploy to all member accounts, the console also creates this role in every member account of your AWS Organization.

Note

The **MGNConnectorInstallerRole** that was previously required to install the connector has been removed and is no longer needed. Its permissions (`mgn:CreateConnector` and `mgn:TagResource`) are now included in the **AWSApplicationMigrationConnectorManagementRole**. When the connector installer runs on the connector machine, it retrieves the **AWSApplicationMigrationConnectorManagementRole** credentials from the AWS Systems Manager agent, which is registered using the SSM hybrid activation, and uses them to register the connector with MGN.

How MGN creates the roles

MGN creates the roles using AWS CloudFormation, with the credentials of the currently signed-in user:

- **Individual account:** MGN creates a CloudFormation stack named **MGNConnectorRoles** in the same AWS Region where you are creating the connector. The stack creates both roles in the current account. IAM roles are global, so the connector can be used in any Region.
- **Multiple accounts (AWS Organizations):** If you are signed in to the organization's management account, or to a delegated administrator account for CloudFormation StackSets, the console displays an **IAM roles deployment scope** option:
 - **This account only** - Create the IAM roles in this account only.
 - **This account and all member accounts** - In addition to the **MGNConnectorRoles** stack in the current account, MGN creates a service-managed CloudFormation StackSet named **MGNConnectorSharingRoles** that deploys the **AWSApplicationMigrationConnectorSharingRole_management-account-id** role to all member accounts in your organization. Automatic deployment is enabled, so accounts that later join the organization automatically receive the role. When you run this from the management account, MGN also enables trusted access between AWS Organizations and CloudFormation StackSets if it is not already enabled.

Note

If deployment to some member accounts fails, role creation in the current account is not rolled back. The console displays a warning with a link to the StackSet operation in the CloudFormation console, where you can review and retry the operation on the failed accounts.

Required permissions

To have MGN create the roles for you, the signed-in user must have permissions to:

- Create CloudFormation stacks with named IAM resources (CAPABILITY_NAMED_IAM), and create IAM roles and attach policies to them.
- For deployment to all member accounts: manage CloudFormation StackSets (create/update StackSets and stack instances), call `organizations:ListRoots`, and, from the management account, enable trusted access for CloudFormation StackSets (`organizations:EnableAWSServiceAccess`).

AWSApplicationMigrationConnectorManagementRole

The **AWSApplicationMigrationConnectorManagementRole** role is the role that is assumed by the connector. The console creates an AWS Systems Manager hybrid activation for this role, and the connector machine registers with Systems Manager as a hybrid managed instance using that activation. From that point, the SSM agent on the connector machine provides the role's credentials to the connector, first to the installer, which registers the connector with MGN (`mgn:CreateConnector`), and then to the connector software for its ongoing operation.

The role is created with the following trust relationship:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "ssm.amazonaws.com"
      },
    },
  ],
}
```



```

        "Action": "sts:AssumeRole"
    }
]
}

```

The following policies are attached to the role:

1. **AmazonSSMManagedInstanceCore** (AWS managed policy)
2. **MgnConnectorPolicy** (inline policy)

The **MgnConnectorPolicy** inline policy contains the following permissions. In the example, replace **111122223333** with the ID of the account in which the role is created, and **us-east-1** with the Region in which you added the connector:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "mgn:CreateConnector",
      "Resource": "arn:aws:mgn:us-east-1:111122223333:*"
    },
    {
      "Effect": "Allow",
      "Action": "mgn:TagResource",
      "Resource": "arn:aws:mgn:us-east-1:111122223333:connector/*",
      "Condition": {
        "StringEquals": {
          "mgn:CreateAction": "CreateConnector"
        }
      }
    },
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Resource": "arn:aws:iam::*:role/AWSApplicationMigrationConnectorSharingRole_111122223333"
    },
    {
      "Effect": "Allow",
      "Action": "secretsmanager:GetSecretValue",
      "Resource": "arn:aws:secretsmanager:us-east-1:111122223333:secret:*"
    }
  ]
}

```

```

        "Condition": {
            "Null": {
                "aws:ResourceTag/AWSApplicationMigrationServiceManaged": "false"
            }
        },
        {
            "Effect": "Allow",
            "Action": "s3:GetObject",
            "Resource": [
                "arn:aws:s3::aws-application-migration-service-us-east-1/latest/source-automation-client/linux/ssaf-client/ssaf_client",
                "arn:aws:s3::amazon-ssm-us-east-1/*"
            ]
        },
        {
            "Effect": "Allow",
            "Action": [
                "logs:CreateLogGroup",
                "logs:CreateLogStream",
                "logs:DescribeLogGroups",
                "logs:DescribeLogStreams",
                "logs:PutLogEvents"
            ],
            "Resource": [
                "arn:aws:logs:us-east-1:111122223333:log-group:/aws/ssm/*",
                "arn:aws:logs:us-east-1:111122223333:log-stream:*"
            ]
        }
    ]
}

```

These permissions allow the connector to:

- Register itself with MGN (`mgn:CreateConnector`) and tag the connector resource during registration (`mgn:TagResource`). These permissions were previously part of the separate **MGNConnectorInstallerRole**, which is no longer used.
- Assume the **AWSApplicationMigrationConnectorSharingRole_ACCOUNT-ID** role in any account in which the role exists, to perform actions on source servers such as installing the AWS Replication Agent.
- Read source server credentials from AWS Secrets Manager. Access is limited to secrets that are tagged with the `AWSApplicationMigrationServiceManaged` tag.

- Download the source server automation client and the AWS Systems Manager agent from the AWS owned Amazon S3 buckets.
- Send connector logs to Amazon CloudWatch Logs, in log groups under the `/aws/ssm/` prefix.

Note

The console-created policy scopes its resources to the connector's account and Region.

AWSApplicationMigrationConnectorSharingRole_management-account-id

The **AWSApplicationMigrationConnectorSharingRole_management-account-id** role is assumed by the **AWSApplicationMigrationConnectorManagementRole** role and by the MGN service in order to perform actions on source servers in the account, for example, to install the AWS Replication Agent. The role name includes the ID of the account that owns the connector (the management account), which allows a single member account to hold sharing roles for multiple management accounts.

MGN supports cross-account scenarios: the role must exist in *every* account that contains source servers on which the connector installs agents. The console creates the role in the current account, and can deploy it to all member accounts of your AWS Organization. To create the role in an account that the console does not deploy to, for example, an account that is not part of your organization, or when you selected **This account only**, download the Member Account Template and deploy it in that account, as described in [Download the CloudFormation templates](#).

The role is created with the following trust relationship, where *management-account-id* is the account in which the connector was added:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "mgn.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
```

```

        "aws:SourceAccount": "111122223333"
      },
      "ArnLike": {
        "aws:SourceArn": "arn:aws:mgn:*:111122223333:*"
      }
    },
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::111122223333:role/
AWSApplicationMigrationConnectorManagementRole"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}

```

The `aws:SourceAccount` and `aws:SourceArn` conditions limit the MGN service to assuming the role only on behalf of resources in the management account, which protects against the confused deputy problem.

The following policy is attached to the role:

1. **AWSApplicationMigrationAgentInstallationPolicy** (AWS managed policy)

Download the CloudFormation templates

If you prefer to review the roles or deploy them in your own workflow, expand **View or deploy roles yourself** in the **IAM Roles** section of the **Add connector** page, and download the CloudFormation templates:

- **Management Account Template** (`mgn-connector-management-roles.json`), creates both the **AWSApplicationMigrationConnectorManagementRole** and the **AWSApplicationMigrationConnectorSharingRole_management-account-id** roles. Deploy this template in the account that owns the connector.
- **Member Account Template** (`mgn-connector-sharing-role.json`), creates only the **AWSApplicationMigrationConnectorSharingRole_management-account-id** role. Deploy this template in every member account whose source servers the connector manages, for example by using a CloudFormation StackSet.

If you deploy the templates yourself, clear the **Create IAM Roles** check box before you continue.

Create roles manually

To create permissions manually, you create the

AWSApplicationMigrationConnectorManagementRole needed to install and run the connector.

The connector assumes the **AWSApplicationMigrationConnectorSharingRole_management-account-id** role as needed, for example, to install the replication agent on a source server.

Note

The **MGNConnectorInstallerRole** is no longer required and does not need to be created. The permissions to register the connector (`mgn:CreateConnector` and `mgn:TagResource`) are included in the **MgnConnectorPolicy** below. The connector installer obtains the **AWSApplicationMigrationConnectorManagementRole** credentials from the AWS Systems Manager agent, which is registered using the SSM hybrid activation.

AWSApplicationMigrationConnectorManagementRole

The **AWSApplicationMigrationConnectorManagementRole** role is the role that is assumed by the Connector. The connector installer uses this role's credentials, provided by the AWS Systems Manager agent, to register the connector with MGN.

To create the role:

1. In the following JSON, replace the example account ID **111122223333** with your account number, and the example Region **us-east-1** with the connector Region. Then, create a policy from the JSON:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": "mgn:CreateConnector",
      "Resource": "arn:aws:mgn:us-east-1:111122223333:*",
      "Effect": "Allow"
    },
    {
      "Action": "mgn:TagResource",
      "Resource": "arn:aws:mgn:us-east-1:111122223333:connector/*",

```

```

        "Effect": "Allow",
        "Condition": {
            "StringEquals": {
                "mgn:CreateAction": "CreateConnector"
            }
        },
        {
            "Action": "sts:AssumeRole",
            "Resource": "arn:aws:iam::*:role/
AWSApplicationMigrationConnectorSharingRole_111122223333",
            "Effect": "Allow"
        },
        {
            "Condition": {
                "Null": {
                    "aws:ResourceTag/AWSApplicationMigrationServiceManaged": "false"
                }
            },
            "Action": "secretsmanager:GetSecretValue",
            "Resource": "arn:aws:secretsmanager:us-east-1:111122223333:secret:*",
            "Effect": "Allow"
        },
        {
            "Action": "s3:GetObject",
            "Resource": [
                "arn:aws:s3:::aws-application-migration-service-us-east-1/latest/
source-automation-client/linux/ssaf-client/ssaf_client",
                "arn:aws:s3:::amazon-ssm-us-east-1/*"
            ],
            "Effect": "Allow"
        }
    ]
}

```

2. If you created an S3 bucket for SSM logging, replace **LOGS-BUCKET** with the bucket name and append the following to the policy:

```

{
    "Action": "s3:PutObject",
    "Resource": "arn:aws:s3:::LOGS-BUCKET/*",
    "Effect": "Allow"
}

```

3. In order for the MGN connector to send logs to CloudWatch, append this statement to the policy:

```
{
  "Effect": "Allow",
  "Action": [
    "logs:CreateLogGroup",
    "logs:CreateLogStream",
    "logs:DescribeLogGroups",
    "logs:DescribeLogStreams",
    "logs:PutLogEvents"
  ],
  "Resource": "*"
}
```

4. Name the policy **MgnConnectorPolicy**
5. Create a role with the following trust relationship:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "ssm.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

6. Attach the following policies:
 - a. **AmazonSSMManagedInstanceCore**
 - b. **MgnConnectorPolicy**
7. Name the role **AWSApplicationMigrationConnectorManagementRole**

AWSApplicationMigrationConnectorSharingRole_management-account-id

The **AWSApplicationMigrationConnectorSharingRole_management-account-id** role is assumed by the **AWSApplicationMigrationConnectorManagementRole** to perform actions on source

servers in member accounts. The role name includes the ID of the account that owns the connector (the management account), which allows a single member account to hold sharing roles for multiple management accounts.

To create the role:

1. Create a policy from the following JSON:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "mgn:StartAgentless",
      "Resource": "arn:aws:mgn:*:*:source-server/*"
    }
  ]
}
```

2. Name the policy **AWSApplicationMigrationAgentInstallationPolicy**.
3. Create a role with the following trust relationship, where *management-account-id* is the account in which the connector was created:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "mgn.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "111122223333"
        },
        "ArnLike": {
          "aws:SourceArn": "arn:aws:mgn:*:111122223333:*"
        }
      }
    }
  ],
  {
```



```

        "Effect": "Allow",
        "Principal": {
            "AWS": "arn:aws:iam::111122223333:role/
AWSApplicationMigrationConnectorManagementRole"
        },
        "Action": "sts:AssumeRole"
    }
]
}

```

4. Attach the **AWSApplicationMigrationAgentInstallationPolicy** policy to the Permission policies.
5. Name the role **AWSApplicationMigrationConnectorSharingRole_management-account-id**, replacing *management-account-id* with the ID of the account in which the connector was created.

Deploy role using a CloudFormation template

For multiple accounts, you can use CloudFormation StackSets to deploy the **AWSApplicationMigrationConnectorSharingRole_management-account-id** role to member accounts. The MGN console provides templates for download, as described in [Download the CloudFormation templates](#) under "Create roles using the MGN console".

If you prefer to deploy the roles using CloudFormation StackSets manually, you can download the templates from the **Add connector** page:

1. On the **MGN connectors** page, choose **Add connector**.
2. Under **IAM Roles**, expand **View or deploy roles yourself** to access the template download options.
3. Download the **Member Account Template** (`mgn-connector-sharing-role.json`) to deploy to member accounts.
4. Using the CloudFormation console or AWS CLI, create a StackSet from the template and deploy it to the desired member accounts. For instructions, see [Getting started with AWS CloudFormation StackSets](#) in the *AWS CloudFormation User Guide*.

Installing the MGN connector on a secured network

The MGN connector and the AWS Replication Agents that the MGN connector installs, require network access to various AWS endpoints. If your on-premises network is not open to AWS

endpoints, then you can install the MGN connector and the AWS Replication Agents with the aid of PrivateLink.

You can connect your on-premises network to your VPCs using AWS VPN or DirectConnect.

Global view

If you are using the [Global view](#) feature, which provides cross-account view and operations, you will have at least one staging VPC per member account.

You will also need to designate a VPC in the management account in order to allow the MGN connector to communicate with AWS services via PrivateLink. If you are migrating some of your source servers into the management account, you can use the same VPC as a staging VPC.

The following sections apply to the MGN connector VPC as well as to each staging VPC.

Create VPC endpoints

To allow the MGN connector and AWS Replication Agents to communicate with AWS services, create the VPC endpoints listed below. For each endpoint:

1. Select your staging area VPC or MGN connector VPC (see [Global view](#) above).
2. **Enable private DNS names.**
3. Choose a subnet, and ensure that a route exists from the MGN connector or AWS Replication Agent to the selected subnet.
4. Ensure that the security groups associated with the endpoint allow inbound traffic from the MGN connector and source servers.

Create the following interface endpoints:

1. **com.amazonaws.*region*.ssm** – The endpoint for the Systems Manager service. This endpoint is required by the SSM Agent, which is installed by the MGN connector installer.
2. **com.amazonaws.*region*.ec2messages** – Systems Manager uses this endpoint to make calls from the SSM Agent to the Systems Manager service.
3. **com.amazonaws.*region*.ssmmessages** – This endpoint is required only if you wish to connect to the MGN connector using Session Manager.
4. **com.amazonaws.*region*.kms** – This endpoint is required only if you wish to connect to the MGN connector using Session Manager and using AWS KMS encryption to add an additional

layer of encryption to the session. For more information, see [Turn on KMS key encryption of session data](#) in the *Amazon Systems Manager User Guide*.

5. **com.amazonaws.*region*.s3** – Systems Manager uses this endpoint to update the SSM Agent and to perform patching operations. The MGN connector installer and the AWS Replication Agent installer download installation assets from this endpoint.
 - a. Note that private DNS names are disabled by default for the S3 endpoint.
 - b. If you wish to also **Enable private DNS only for inbound endpoint**, you must first create an S3 gateway VPC endpoint. For more information, see [S3 Private DNS](#) in the *Amazon Simple Storage Service User Guide*.
6. **com.amazonaws.*region*.secretsmanager** – The MGN connector calls this endpoint to retrieve source server credentials.
7. **com.amazonaws.*region*.sts** – The MGN connector calls this endpoint to retrieve credentials of the AWS Replication Agent installer role.
8. **com.amazonaws.*region*.mgn** – The endpoint for MGN. This endpoint is required by the MGN connector, the AWS Replication Agent, and their respective installers. If a VPCE Policy is used (to scope down access), add the following statement to your policy:

```
{
  "Effect": "Allow",
  "Principal": "*",
  "Action": "execute-api:Invoke",
  "Resource": "arn:aws:execute-api:<region>:*:*/POST/CreateSessionForMgn"
}
```

For more information, see [Creating an interface endpoint](#) in the *Amazon VPC User Guide*.

Create a Route 53 inbound endpoint

To route your traffic to the VPC endpoints created above, create a Route 53 inbound endpoint in your staging area VPC or the MGN connector VPC (see [Global view](#) above).

Ensure that the security group associated with the inbound endpoint allows traffic from your on-premises DNS resolvers.

Configure DNS resolvers on your on-premises network to forward DNS queries for the endpoints of the above AWS services, to the IP addresses of your Route 53 inbound endpoint. To find the regional endpoints of these services, see [Service endpoints](#) in the *AWS General Reference Guide*.

For example, the endpoint of the MGN service in the US East (Ohio) Region (us-east-2) is **mgn.us-east-2.amazonaws.com**

For more information, see [Forwarding inbound DNS queries to your VPCs](#) in the *Amazon Route 53 User Guide*.

Modify replication settings

In order to allow the AWS Replication Agent to communicate with the replication server without using the public internet, you must use Private IP for data replication. The replication server requires access to the EC2 service. Therefore:

- If your staging area VPC has a VPC endpoint for **com.amazonaws.*region*.ec2** with private DNS names enabled, or if your staging area subnet has a route to the public internet via a NAT gateway, then the replication server can communicate with EC2 over its private IP. Choose the option:

Use private IP for data replication

- Otherwise, if your staging area subnet has a route to the public internet via an internet gateway, a public IP is required for the replication server to reach EC2. Choose the option:

Create public IP, and use Private IP for data replication

Ensure that the security groups associated with the MGN VPC endpoint allow inbound traffic from the replication server.

Verify VPC endpoints are being used

Use CloudTrail to verify that calls to AWS services from the MGN connector and its associated source servers, are made via the **vpcEndpointIds** of the VPC endpoints you have created.

Review details about your MGN connectors

Click the **MGN connector name** of any MGN connector to open its details page.

The page includes the following details:

- Overview - View all the information related to the specific MGN connector including state and when it last communicated with AWS Transform MGN.

- **Source servers** - This section features all the source servers managed by the MGN connector.

Each row in the "servers" table provides details about a single source server, including:

- **Hostname** - The source server's hostname.
- **Account** - The source server account id.
- **Prerequisites** - The status of prerequisites verification, with the following options: **Verified**, **Not verified**, **In progress**, or **Invalid**.
- **Agent installed** - Indicates whether the AWS MGN Agent is installed on the server.
- **Credential secret** - The secret of the specific source server.
- **Next step** - What is the new action in the connector installation workflow. Options include:
 - **Initiate test** - Test your source server before migration.
 - **Mark as tested** - Mark that the source server is ready for migration.
 - **Check prerequisites** - Ensure that the source server meets the required prerequisites for replication agent installation.
 - **Wait for check to complete** - This indicates that the prerequisites for replication agent installation are being verified. If this step completes successfully, the next step is **Install agent**. If not, the next step is **Resolve cause of invalidity**.
 - **Resolve cause of invalidity** - This indicates that the prerequisite verification process failed and that a specific issue needs to be resolved.
- **Tags** - This section features the tags associated with your connector.

Manage your MGN Connectors

The MGN Connectors page lists all the installed MGN connectors, providing a quick overview of your MGN connectors and their status and allowing you to quickly perform actions.

Introduction to the MGN connector page

The **MGN Connectors** page displays the list of MGN connectors, and supports adding, deleting and editing MGN connectors as well as performing actions using the MGN connectors.

The **MGN Connectors** page provides information for each MGN connector, including:

- **MGN Connector name** - The unique name for each MGN connector. Additional details of the MGN connector are available in the MGN details page. Click the MGN connector name, to view its details.
- **Registered servers** - The number of registered source servers managed by this MGN connector.
- **Last seen** - The last time AWS Transform MGN communicated with the MGN connector.

Topics

- [Edit connector](#)
- [Delete MGN connector](#)
- [Register server credentials](#)
- [Verify source server prerequisites](#)
- [Install the replication agent](#)
- [View command history](#)

Edit connector

To edit an MGN connector, click **Edit**. When the **Edit MGN connector** page opens, you can modify the MGN connector name and tags.

To finalize your changes, click **Save changes**.

Delete MGN connector

To delete an MGN connector, click **Delete**. When the **Delete MGN connector** dialog opens, verify that you want to delete the selected MGN connector. Once an MGN connector is uninstalled, it can no longer be used to manage your source servers.

Note

Deleting the MGN connector will disassociate the servers from the MGN connector, but will not delete them from servers inventory.

Register server credentials

Once you have the MGN connector set up and ready to use, you can register source servers to the MGN connector. To do so click on the MGN connector name, then click "Register servers".

The servers list contains the source servers that were imported via the import feature or discovered by the agentless replication process.

Select the source servers you want to register to the MGN connector. Click the "Register servers with the MGN connector" button.

To perform actions on your source server, you must provide source server credentials. Server credentials are stored in AWS Secrets Manager. You can use an existing secret from AWS Secrets Manager or create a new one. You can create the credentials in the MGN console by choosing **Register server credentials** from the **Actions** menu.

- Use existing secret
 - Using AWS Secrets Manager, MGN can use the stored source server credentials to connect to the source machine and perform actions on it. You must specify the secret that stores the source server credentials.
 - You may designate the same secret for multiple source servers if they share the same credentials.
 - Be sure to add the `AWSApplicationMigrationServiceManaged` tag to the secret with the value set to `True`.
- Create new secret
 - **Secret name** - Enter a name for your new secret. The name you specify will be saved in AWS Secrets Manager.
 - **Encryption key** - To encrypt, either use the KMS key provided by Secrets Manager or create your own customer managed KMS key.
 - **For Windows servers:**
 - **Communication protocol** - This is the WinRM connection protocol between the MGN Connector and source servers used to install the agents. Specify either HTTP or HTTPS. We recommend HTTPS for secure and encrypted communication. Default is HTTPS.
 - **UserName** - A user that is authorized to install the agent and perform actions on the source server.
 - **Password** - The specific source server's password.

- **CertificateAuthority** (Required if WinCaValidation is true) - The CA public certificate in PEM format, base64-encoded. Must be omitted or empty if WinCaValidation is false.
- **For Linux servers:**
 - **UserName** - A user that is authorized to install the agent and perform actions on the source server.
 - **Credentials** - Provide one of the following:
 - **PrivateKey** - The source server's RSA private key in PEM format, base64-encoded. (The connector uses RSA keys only.)
 - **Password** - The specific source server's password (alternative to PrivateKey).
 - **HostKey** (Required if LinuxHostKeyValidation is true) - The source server's public host key in the format: algorithm_name base64_public_key (e.g., ssh-ed25519 AAAA...). This is the full base64 public key blob from ssh-keyscan output, not a fingerprint. Must be omitted if LinuxHostKeyValidation is false.
- **Tags** - Secret key-value pairs will be assigned to the new secret. Note that AWSApplicationMigrationServiceManaged tag will also be added with the value set to True.
- Here is the structure of the secrets manager entry:

```
{
  "WinConnectionProtocol": "HTTPS",
  "WinUserName": "windows_username",
  "WinPassword": "windows_password",
  "WinCertificateAuthority": "<base64-encoded CA public certificate (PEM)>",
  "WinCaValidation": true,
  "LinuxUserName": "linux_username",
  "LinuxPrivateKey": "<base64-encoded RSA private key (PEM)>",
  "LinuxPassword": "<alternative to LinuxPrivateKey>",
  "LinuxHostKey": "ssh-ed25519 <base64 public host key>",
  "LinuxHostKeyValidation": true
}
```

- **Important notes on secret format:**
 - **Base64 encoding:** LinuxPrivateKey and WinCertificateAuthority must be base64-encoded PEM values. If you create the secret in the MGN console, encoding is performed

automatically. If you create the secret manually in AWS Secrets Manager, you must encode these values yourself. Raw PEM text will cause connection failures.

- **Private key format:** `LinuxPrivateKey` must be an RSA private key. Other key types are not supported.
- **Validation flags:** `WinCaValidation` and `LinuxHostKeyValidation` are required fields and must always be present in the secret.
- **Validation constraints:** `WinCertificateAuthority` is required only when `WinCaValidation` is true. `LinuxHostKey` is required only when `LinuxHostKeyValidation` is true. When their validation flag is false, these fields should be omitted or empty.
- **HTTP protocol constraint:** `WinCaValidation` must be false when `WinConnectionProtocol` is set to HTTP.
- **HostKey format:** The `LinuxHostKey` value is the full base64 public key blob from `ssh-keyscan` output (the second field), in the format `algorithm_name base64_key`. It is NOT a fingerprint or thumbprint.

 Note

The CA/HostKey validation is controlled by the validation flags (`WinCaValidation` and `LinuxHostKeyValidation`). When a validation flag is set to true, you must provide the corresponding CA or HostKey value. When set to false, the CA or HostKey fields should be omitted or left empty. If you do not provide the required value when validation is enabled, credential validation will fail with a "mandatory field not found" error.

Verify source server prerequisites

The **Verify prerequisites** action ensures the AWS replication agent can be installed on each of the source servers. The verification process ensures there's enough disk space, RAM and CPU for installing the AWS replication agent.

Install the replication agent

Following the prerequisite check, you can proceed to **install the replication agent**, to start your migration execution.

View command history

After performing an action, you can **view the command history** for information on the command status.

Import and Export a Data Inventory to Coordinate a Migration

The import and export feature allows you to easily plan and coordinate your migrations.

Use this feature to import and export your source servers, applications, and waves from and to a CSV file. The file can also include launch template attributes to simplify bulk configurations.

- Import from a CSV file – Import your data from a local disk or an S3 bucket and create entities in your account.
- Export to a CSV file – Export data from your account to a local disk or an S3 bucket and merge it into a single file that you can easily review and process offline.

Topics

- [Editing your configuration](#)
- [Importing your data inventory](#)
- [Exporting your data inventory](#)
- [View export history](#)

Editing your configuration

You can use the **Import and Export** feature to edit your server configuration, including bulk changes:

1. Export the data to a CSV file.
2. Edit your file.
3. Import the edited file back into the service.

Note

When editing a CSV file in Microsoft Excel be sure to save it in the **MS-DOS** CSV data format.

When processing the inventory file import, values provided in the file overwrite existing values. *Cells that are empty are ignored.* This allows you, when editing the file, to include values only in the inventory columns that you want to overwrite.

Deleting values from the inventory file

To delete a value replace the existing value with the string `[RESET_VALUE]` and then save and upload the file. This clears the value in the MGN resource or the Amazon EC2 launch template.

Using `[RESET_VALUE]` is available for these parameters:

- `mgn:launch:iam-instance-profile:name`
- `mgn:launch:nic:0:network-interface-id`
- `mgn:launch:nic:0:private-ip:0`
- `mgn:launch:nic:0:security-group-id:0`
- `mgn:launch:nic:0:subnet-id`
- `mgn:launch:tag:instance:key1`
- `mgn:launch:volume:/dev/sda:type`

Deleting tags from the inventory file

You can also delete tag fields. To delete the tag value replace the existing value with the string `[RESET_VALUE]`. To delete both the key and the value replace the existing value with the string `[DELETE_TAG]`.

Note

Tags that are created by the service, such as `AWSApplicationMigrationServiceManaged`, can't be changed or deleted. Attempting to do so results in an error.

Importing your data inventory

The **Import** feature allows you to easily import your inventory of servers, applications, and waves from a CSV file that is saved in your local disk or an S3 bucket.

Note

The import feature is not supported for IPv6.

Topics

- [Define required permissions for importing](#)
- [Inventory Import parameters](#)
- [Importing your data inventory from a local disk](#)
- [Importing your data inventory from an S3 bucket](#)
- [View import history](#)

Define required permissions for importing

In order to use the import feature, you will need to create a role with the following policies (or any extension of them):

Managed policies:

- [AWSApplicationMigrationFullAccess](#)
- [AWSApplicationMigrationEC2Access](#)

Note

These managed policies grant broad access to MGN features beyond the import function. For least privilege, use IAM Access Analyzer to generate a policy scoped to only the actions required for your import workflow. For more information, see [Generate policies based on access activity](#).

Additional policies:

```
{
  "Sid": "AllowS3Access",
  "Effect": "Allow",
  "Action": [
    "s3:GetObject"
```

```
],  
  "Resource": "arn:aws:s3:::amzn-s3-demo-bucket/*"  
}
```

When starting an import on an Amazon S3 bucket source that is owned by another account, ensure that the role or user has access to the Amazon S3 objects. When using the API, the Amazon S3 bucket owner parameter defaults to the current user's account ID.

The following is an example of an S3 bucket policy in the target account:

JSON

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Sid": "ExampleStatement",  
      "Effect": "Allow",  
      "Principal": {  
        "AWS": "arn:aws:iam::123456789012:user/Dave"  
      },  
      "Action": [  
        "s3:GetObject"  
      ],  
      "Resource": "arn:aws:s3:::amzn-s3-demo-bucket/*"  
    }  
  ]  
}
```

Note

If the Amazon S3 objects are encrypted with SSE-KMS, ensure that the role or user initiating the import has access to decrypt using the AWS KMS key. This feature does not support SSE-C encrypted Amazon S3 objects.

Inventory Import parameters

The imported file can include multiple parameters, including:

Parameter	Description
mgn:account-id	The ID of the account into which to import. This account must be managed by the calling account. Defaults to the calling account.
mgn:app:description	The description of the application being imported.
mgn:app:id	The ID of the application being imported. Use this parameter to update an existing application.
mgn:app:name	The name of the application being imported.
mgn:app:tag:appkey1	The value of the application tag key (in this example, the tag key is appkey1).
mgn:launch:boot-mode	The boot mode of the launch instance. Expected values: LEGACY_BIOS, UEFI, or USE_SOURCE.
mgn:launch:copy-private-ip	The option to copy the private IP of the source server to the launch instance (true or false).
mgn:launch:iam-instance-profile:name	The name of the instance profile associated with the launch instance.
mgn:launch:instance-type	The EC2 instance type of the launch instance (for example, m4.large).
mgn:launch:map-tag-key	The key of the map migration tag applied by map auto tagging. Expected values: map-migrated or aws-apn-id. Values are case-sensitive.
mgn:launch:map-tag-value	The value of the map migration tag applied by map auto tagging, for example, your MAP program engagement (MPE) ID.

mgn:launch:map-tagging	The option to enable map auto tagging, which automatically applies the map migration tag to resources launched by MGN (true or false).
mgn:launch:nic:0:network-interface-id	The ID of the network interface that appears first in the launch template ("0" refers to the first network interface, "1" would refer to the second network interface, and so on).
mgn:launch:nic:0:private-ip:0	The private IP that appears first in the network interface that appears first in the launch template.
mgn:launch:nic:0:security-group-id:0	The security group that appears first in the network interface that appears first in the launch template.
mgn:launch:nic:0:subnet-id	The subnet ID that appears first in the network interface that appears first in the launch template.
mgn:launch:operating-system-licensing	The operating system licensing approach. Expected values: BYOL (bring your own license) or LI (license included). Values are case-sensitive and must be uppercase.
mgn:launch:placement:host-id	The host ID of the placement of the launch instance.
mgn:launch:placement:tenancy	The tenancy of the launch instance. Expected values: default, dedicated, or host. Values are case-sensitive and must be lowercase.
mgn:launch:post-actions:enabled	The option to enable post-launch actions for the server (true or false).
mgn:launch:post-actions:<action-name>:ssmDocumentName	The name of the SSM document that the post-launch action runs (in this example, the action name is <action-name>).

mgn:launch:post-actions:<action-name>:order	The order in which the post-launch action runs, relative to other post-launch actions. Must be a positive integer.
mgn:launch:post-actions:<action-name>:active	The option to activate the post-launch action (true or false).
mgn:launch:post-actions:<action-name>:mustSucceedForCutover	The option that determines whether the post-launch action must complete successfully before cutover can proceed (true or false).
mgn:launch:post-actions:<action-name>:timeoutSeconds	The timeout of the post-launch action, in seconds. Must be an integer greater than 0.
mgn:launch:post-actions:<action-name>:description	The description of the post-launch action.
mgn:launch:post-actions:<action-name>:parameters	The parameters passed to the SSM document of the post-launch action, as a JSON object with two optional members, <code>parameters</code> and <code>externalParameters</code> . No other members are allowed. <code>parameters</code> maps each document parameter name to an ordered list of <code>{"value": "<literal>", "type": "String"}</code> entries, where type is <code>String</code> (default) or <code>StringList</code> . <code>externalParameters</code> maps each document parameter name to a dynamic path string. For example: <code>{"parameters":{"Region":[{"value":"us-east-1"}]},"externalParameters":{"InstanceId":"\$.instanceId"}}</code>
mgn:launch:start-instance	The option to automatically start the test or cutover instance after launch (true or false). If set to false, the instance is launched in a stopped state.

mgn:launch:tag:instance:key1	The value of launch instance tag "key1" (in this example, the tag key is key1).
mgn:launch:volume:/dev/sda:type	The EBS volume type of the launch instance's volume whose device name is /dev/sda. The device name must match the name in the EC2 launch template's block device mappings (for example, /dev/sda1 or xvdf).
mgn:region	The AWS Region to which you are importing , which must be the Region of your MGN console. If left blank, defaults to the console Region.
mgn:replication:associate-default-security-group	The option to associate the default MGN security group with the replication servers (true or false).
mgn:replication:bandwidth-throttling	The bandwidth throttling limit for data replication, in Mbps. A value of 0 means no throttling.
mgn:replication:create-public-ip	The option to create a public IP address for the replication servers (true or false).
mgn:replication:data-replication-routing	Determines whether data is replicated over private connectivity or over the public internet. Expected values: PRIVATE_IP or PUBLIC_IP.
mgn:replication:default-large-staging-disk-type	The default EBS volume type to use for large staging area disks. Expected values: GP2, GP3, or ST1.
mgn:replication:ebs-encryption	The EBS encryption option for the replicated volumes. Expected values: DEFAULT or CUSTOM.

mgn:replication:ebs-encryption-key-arn	The ARN of the KMS key used for EBS encryption, when using a custom encryption key.
mgn:replication:fsx-ontap:credentials-secret-arn	The ARN of the AWS Secrets Manager secret that holds the credentials of the Amazon FSx for NetApp ONTAP file system used as the staging area. Required when mgn:replication:storage-type is FSX_ONTAP, and rejected when it is EBS.
mgn:replication:fsx-ontap:storage-virtual-machine-id	The ID of the Amazon FSx for NetApp ONTAP storage virtual machine (SVM) used as the staging area. Required when mgn:replication:storage-type is FSX_ONTAP, and rejected when it is EBS.
mgn:replication:internet-protocol	The internet protocol used for data replication. Expected values: IPV4 or IPV6.
mgn:replication:replication-server-instance-type	The EC2 instance type of the replication servers (for example, t3.small).
mgn:replication:security-group-id:0	The security group that appears first in the replication settings ("0" refers to the first security group, "1" would refer to the second security group, and so on).
mgn:replication:staging-area-subnet-id	The ID of the subnet in which the staging area is created.
mgn:replication:staging-area-tag:stagingkey1	The value of the staging area resource tag key (in this example, the tag key is stagingkey1).

mgn:replication:storage-type	The staging area storage type. Expected values: EBS or FSX_ONTAP. When set to FSX_ONTAP, you must also provide mgn:replication:fsx-ontap:storage-virtual-machine-id and mgn:replication:fsx-ontap:credentials-secret-arn.
mgn:replication:store-snapshot-on-local-zone	The option to store snapshots on the Local Zone (true or false).
mgn:replication:use-dedicated-replication-server	The option to use a dedicated replication server for each source server (true or false).
mgn:replication:use-fips-endpoint	The option to use FIPS-compliant service endpoints for data replication (true or false).
mgn:server:fqdn-for-action-framework	The FQDN that the MGN connector uses to connect to the server.
mgn:server:id	The ID of the source server being imported. Use this parameter to update an existing source server.
mgn:server:platform	The server's platform (Linux or Windows).
mgn:server:tag:serverkey1	The value of the server tag key (in this example, the tag key is serverkey1).

mgn:server:user-provided-id	The server's user-provided ID. This parameter is used by MGN to consistently recognize the server replication, and avoid duplication when importing inventory from a CSV file. This parameter should be provided during the AWS replication agent installation by the MGN connector or other installation method. This attribute can be edited via the edit server option. After a replication agent is installed on the server, mgn:server:user-provided-id and mgn:server:fqdn-for-action-framework are the only server parameters that import can still change. Any other server parameter in the row causes the row to fail.
mgn:wave:description	The description of the imported wave.
mgn:wave:id	The ID of the wave being imported. Use this parameter to update an existing wave.
mgn:wave:name	The name of the imported wave.
mgn:wave:tag:wavekey1	The value of the wave tag key (in this example, the tag key is wavekey1).
mgn:launch:transfer-server-tags	The option to transfer any user-configured custom tags from your source servers onto your test or cutover instance.

Additional considerations

Note the following considerations regarding the import parameters:

1. Server entries must include either the server IP address, or the FQDN.
2. If a row provides a property of a resource (e.g. mgn:wave:description is a property of a wave), then that row should also provide a parameter that identifies that resource (as explained in the following considerations).

3. If a resource's ID (mgn:server:id, mgn:app:id, or mgn:wave:id) is provided, the service will look for this resource in order to update it. If this resource is not found, the import will fail.
4. If a resource's ID is not provided, the service will look for a resource's alternative identification:
 - For an application: mgn:app:name
 - For a wave: mgn:wave:name
 - For a server: mgn:server:user-provided-id
5. If a resource's alternative identification exists in MGN, the service will update this resource with new values (if applicable).
6. If a resource's alternative identification does not exist in MGN, the service will create the resource.
7. Two rows that refer to the same resource need not provide the same parameters for that resource, but they must not conflict. For example, if two rows provide the same mgn:wave:name, it is acceptable for one row to provide mgn:wave:description and for the other row to leave the value blank. However, the two rows must not provide conflicting values of mgn:wave:description.

Importing your data inventory from a local disk

To import your inventory from a local disk, take the following steps:

1. Select **Import** from the left-hand navigation menu (under **Import and export**) and you'll be navigated to the **Import inventory** tab.
2. Select **Import from local disk**.
3. Choose **Choose file** to select the CSV file from which you want to import the data.
4. Choose **Import**.

Note

The file will also be automatically imported to an S3 bucket created by MGN. It is highly recommended that you [apply Amazon S3 bucket security practices](#) where your CSV files are stored.

Importing your data inventory from an S3 bucket

To import your inventory from an S3 bucket, take the following steps:

1. Select **Import** from the left-hand navigation menu (under **Import and export**) and you'll be navigated to the **Import inventory** tab.
2. Select **Import from S3**.
3. Choose **Browse** to choose the Amazon S3 storage source from which you want to import the data.
4. Choose **Import**.

Note

It is highly recommended that you [apply Amazon S3 bucket security practices](#) where your CSV files are stored.

[Learn more about S3 permissions and policies.](#)

View import history

Select the **Import history** tab to view the files imported in the last 7 days, including their status and the task's start and end time.

You can change the settings according to your preferences by choosing the settings icon located in the right-hand corner of the screen.

To see all the related task messages, choose the task ID. To copy the messages, choose **Copy**.

Exporting your data inventory

The **Export** feature allows you to easily export your inventory of servers, applications, and waves to a CSV file that is saved in your local disk or an S3 bucket.

Note

The export feature is not supported for IPv6.

Defining required permissions for export

In order to use the export feature, you will need to create a role with the following policies (or any extension of them):

Managed policies:

- `AWSApplicationMigrationReadOnlyAccess`

Additional policies:

```
{
  "Sid": "AllowS3Access",
  "Effect": "Allow",
  "Action": [
    "s3:PutObject"
  ],
  "Resource": "arn:aws:s3:::amzn-s3-demo-bucket/*"
},
{
  "Sid": "AllowMgnStartExport",
  "Effect": "Allow",
  "Action": [
    "mgn:StartExport"
  ],
  "Resource": "*"
}
```

When starting an export on an Amazon S3 bucket source that is owned by another account, ensure that the role or user has access to the Amazon S3 objects. When using the API, the Amazon S3 bucket owner parameter defaults to the current user's account ID.

The following is an example of an Amazon S3 bucket policy in the target account:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ExampleStatement",
```



```
"Effect": "Allow",
"Principal": {
  "AWS": "arn:aws:iam::123456789012:user/Dave"
},
"Action": [
  "s3:PutObject"
],
"Resource": "arn:aws:s3:::amzn-s3-demo-bucket/*"
}
]
```

Note

If the Amazon S3 objects are encrypted with SSE-KMS, ensure that the role or user initiating the export has access to decrypt using the AWS KMS key. This feature does not support SSE-C encrypted Amazon S3 objects.

Required Amazon S3 bucket permissions

Before you create an export job, you must create the destination S3 bucket to export to. AWS Transform MGN doesn't create the S3 bucket for you. The S3 bucket that you specify can't be publicly accessible, and can't be configured as a [Requester Pays](#) bucket. After you create the S3 bucket, confirm that the bucket has the required permissions policy to allow AWS Transform MGN to write the export files to it.

Export parameters

The exported file includes the same parameters as the imported file. For the full list of parameters and their descriptions, see [Import parameters](#).

In addition, the exported file includes the following view-only parameters. These parameters are exported for informational purposes only, and are ignored if they are present in an imported file:

Parameter	Description
mgn:server:lifecycle-state	The server's lifecycle state.

mgn:server:replication-state	The state of the replication.
mgn:server:replication-type	The type of the replication (agent-based or agentless).

Note

If the bucket you're exporting to is encrypted with customer managed keys (KMS), that KMS key's policies must give MGN permission to use it. This permission is given through the user or role that initiates the export job.

If you choose to encrypt your export using a key protected by AWS Key Management Service (AWS KMS), the key must be in the same Region as the destination S3 bucket.

Exporting your data inventory to a local disk

To export your inventory to a local disk, take the following steps:

1. Select **Export** from the left-hand navigation menu (under **Import and export**) and you'll be navigated to the **Export inventory** tab.
2. Select **Export to a local disk**.
3. Specify the name of the CSV file into which you want to download the data.

Note

- The file will also be automatically downloaded to an S3 bucket created by AWS Transform MGN.
- You must have the required permissions to perform this action.
- It is highly recommended that you [apply S3 bucket security practices](#) where your CSV files are stored.

4. Choose **Export**.

Exporting your data inventory to an S3 bucket

To export your inventory to an S3 bucket, take the following steps:

1. Select **Export** from the left-hand navigation menu (under **Import and export**) and you'll be navigated to the **Export inventory** tab.
2. Select **Export to S3 bucket**.
3. Choose **Browse S3** to choose the Amazon S3 storage target to which you want to export the data.
4. Specify the Amazon S3 bucket owner (the current AWS account or a different one) according to your preferences. If you select a different AWS account, you must enter the bucket owner's account ID.

 Note

- You must have write privileges to export an inventory to a specific bucket.
- It is highly recommended that you [apply S3 bucket security practices](#) where your CSV files are stored. [Learn more about S3 permissions and policies.](#)

5. Choose **Export**.

View export history

Select the **Export history** tab to view the files exported in the last 7 days, including their status, progress, and the task's start and end time.

You can change the settings according to your preferences by choosing the settings icon located in the right-hand corner of the screen.

To see all the related task messages, choose the task ID. To copy the messages, choose **Copy**.

Manage large-scale migrations with global view

The AWS Transform MGN global view feature enables you to manage large-scale migrations across multiple accounts. Global view provides visibility, and the ability to perform actions on source servers, apps, and waves in different AWS accounts.

Global view uses AWS Organizations to structure a management account that has access to source servers in multiple member accounts, and member accounts that only have access to their own source servers.

To use this feature:

- You need to have an AWS account in which AWS Transform MGN is initialized.
- The account must be a management account in AWS Organizations, or a delegated admin for AWS Transform MGN which has the same feature permissions as a management account in AWS Organizations.

Setting up your AWS Organization

The AWS Organizations service enables you to consolidate multiple AWS accounts into a single organization that you create and manage. You can create member accounts or invite existing accounts to join your organization. [Learn more about AWS Organizations.](#)

To use global view, first create your organization in the AWS Organizations console:

1. Go to the AWS Organizations console.
2. [Create a new AWS organization.](#)
3. [Invite member accounts](#) you want to manage within MGN.

Activate trusted access for AWS Transform MGN

To use global view, you must activate trusted access to AWS Transform MGN for your organization.

Attach the [AWSOrganizationsFullAccess](#) managed policy to the user.

To enable service access for your organization, take the following steps:

1. Activate trusted access for MGN
 - a. Log in as management account.
 - b. Select **Global view** from the left-hand navigation menu.
 - c. Activate service access by choosing the 'Enable AWS Organizations service access' button

[Learn more about activating trusted access.](#)

2. Select members and turn them into delegated admins for MGN by calling the [RegisterDelegatedAdministrator](#) API, including the service name:

```
{
  "AccountId": "string",
  "ServicePrincipal": "mgn.amazonaws.com"
}
```

Important

You can register up to 5 delegated administrators.

Setting up CloudFormation StackSets

After you set up your organization, you need to configure CloudFormation StackSets to create the required role per management account:

AWSApplicationMigrationSharingRole_<MANAGEMENT_ACCOUNT_ID>.

AWS CloudFormation StackSets extends the capability of stacks by enabling you to create, update, or delete stacks across multiple accounts and AWS Regions with a single operation.

[Learn more about CloudFormation StackSets.](#)

Important

StackSet automatically creates the roles in all accounts. You can choose to create the roles manually in each member account of the organization, however, this must be done for each account individually.

To set up your StackSet:

1. Go to the CloudFormation console.
2. Select **StackSets**.
3. Choose the **Activate trusted access** button.
4. Create StackSet.
5. On the **Choose a template** page, under **Prerequisites – prepare template**, choose **Use a sample template**.
6. Under **Select a sample template**, select **Create roles to access multiple accounts via AWS Transform MGN**, and choose **Next**.
7. Provide the name and description or use the existing values.
8. Under **Parameters**, add the account ID of each admin or delegated admin and choose **Next**.
9. Select or provide the required parameters .

 Important

- Under **Deployment targets**, select **Deploy to organization**.
- Select only one specific AWS Region – we recommend that you select your StackSet Region.
- To provide enhanced stability, we recommend that you set the **Failure tolerance optional** to a high value - at least as high as the number of accounts within the organization.

10. Check the box next to **I acknowledge that AWS CloudFormation might create IAM resources with custom names** and choose **Submit**.

Once all the steps are completed, you should be able to see your new StackSet in **StackSet details > Stack instances**.

Using an AWS KMS customer managed key for encryption in member account

If you decide to use a customer managed key, or if your default Amazon EBS encryption key is a customer managed key in member account, you must add permissions to the

AWSApplicationMigrationSharingRole_<MANAGEMENT_ACCOUNT_ID> to allow management account to use it.

Using Administrator access, add these permissions to the AWSApplicationMigrationSharingRole_<MANAGEMENT_ACCOUNT_ID>:

Inviting an AWS account to join your organization

After you create an organization and verify the email address associated with the management account, you can invite existing AWS accounts to join your organization. Only management accounts can send an invitation to other accounts.

When you invite an account, AWS Organizations sends an invitation to the account owner, who decides whether to accept or decline the invitation. You can use the AWS Organizations console to initiate and manage invitations that you send to other accounts.

[Learn how to send invitations to other AWS accounts.](#)

Using global view

Use the global view feature to see source servers across various member accounts and to perform various actions such as installing the SSM Agent.

To use global view attach the [AWSOrganizationsReadOnlyAccess](#) managed policy to the user.

The main **Global view** page provides an overview of your account. The information differs for a management account and a member account.

- Management account: Displays **Account information** that includes the AWS organizations permissions, number of linked accounts, and the total number of source servers, applications, and waves. The **Linked account** section displays the relevant information only for the linked accounts.
- Member account: Displays the **Account information** that includes the AWS organizations permissions, and the number of source servers, applications, and waves in the specific account.

As a management account, you are able to choose **All accounts** and **My account** from the drop-down menu, changing your view of source servers, applications, or waves.

Source servers in member accounts

As a management account, you can view source servers in your account and all member accounts. You can also perform specific actions on managed servers.

Single managed source server

As a management account, you can perform the following actions on a single managed source server.

- Change staging disk type
- Edit replication settings
- Launch settings – edit general launch settings only
- Post launch
 - Deactivate the post-launch feature for this server
 - Change deployment settings (test and cutover, test only, or cutover only)
- Start/stop replication
- Test and cutover drop-down menu:
 - Launch test
 - Mark as ready for cutover
 - Revert to ready for testing
 - Launch cutover
 - Finalize cutover
 - Revert to ready for cutover
 - Terminate launch instances

Multiple managed source server

As a management account, you can perform the following actions on multiple managed source servers.

- Edit replication settings – the edited servers must be from the same account
- Add server to application – the added servers must be from the same account
- Disconnect servers from service

- Mark as archived
- Start/stop replication

Applications

As a management account, you can perform the following actions on a single or multiple managed applications:

- Add application
- Edit application
- Delete application
- Test and cutover drop-down menu (these actions can also be performed on multiple applications):
 - Launch test
 - Mark as ready for cutover
 - Revert to ready for testing
 - Launch cutover
 - Finalize cutover
 - Revert to ready for cutover
 - Add application to wave
 - Start/stop replication
 - Archive application

Waves

As a management account, you can perform the following actions on a single managed wave:

- Add wave
- Edit wave
- Delete wave
- Test and cutover drop-down menu (these actions can also be performed on multiple waves):
 - Launch test
 - Mark as ready for cutover

- Revert to ready for testing
- Launch cutover
- Finalize cutover
- Revert to ready for cutover
- Add application to wave
- Start/stop replication
- Archive application

Import/Export

Use this feature to import and export your source servers, applications, and waves from a single or multiple accounts using the CSV template file.

Launching test and cutover instances

AWS Transform MGN allows you to launch test and cutover instances in AWS. Before launching instances, you must configure your Launch settings. The following documentation explains how to configure Launch settings and how to launch Test and cutover instances using the configured settings.

Launch settings determine how your test and cutover instances are launched in AWS. Through Launch settings, you can fully customize your test and cutover instances by configuring key metrics, such as the subnet within which the instance is launched, the instance type to be used, licence transfers, replication status, and a variety of other settings. MGN ensures that your test and cutover instances constantly abide by the latest AWS security, instance, and other updates by using EC2 launch templates. EC2 launch templates always use the latest EC2 instance and technology. They integrate with AWS Transform MGN to give you full control over every single setting within your test and cutover instance. Once you have configured your instance's launch settings, you can launch them directly through the MGN console. During the launch process, either during test or cutover instance launch, the AWS replication agent is removed from the test or cutover instance, and does not run on it.

Topics

- [Preparing for test and cutover instance launch](#)
- [Launch settings](#)
- [Launching test instances](#)
- [Launching cutover instances](#)
- [Review launch history](#)

Preparing for test and cutover instance launch

Before launching your instances, you must ensure that your environment is set up properly to ensure successful launches. Check the following before continuing:

- Prepare your subnets for launch - Plan which subnets you will use to launch your test and cutover instances. You use these subnets in your EC2 launch template when you configure your Launch settings.

- Create security groups within the subnets - Create the Security groups you want to use within your prepared subnets. You set these Security groups in your EC2 Launch template when you configure launch settings.

Note

Customers that want to run a proof of concept can skip this step. AWS Transform MGN automatically uses the default subnet and Security groups. Make sure that you have not deleted your default subnet.

Launch settings

The launch settings include two sections: the general launch settings, and the EC2 launch template, which determine how a test or cutover instance is launched for each source server in AWS.

Launch settings, including the EC2 launch template, are automatically created each time you add a source server to AWS Transform MGN.

The launch settings can be modified at any time, including before the source server has completed its initial sync.

Note

Any changes made to the launch settings only affect newly launched test and cutover instances.

Note

For many customers, there is no need to modify the launch settings or the EC2 launch template to launch test or cutover instances.

Launch settings can only be changed for one server at a time through the AWS Transform MGN console.

Note

You can modify launch settings for multiple servers at a time by using the AWS Transform MGN API.

You can access the launch settings of a specific source server through the server details view by choosing its Hostname from the **Source servers** page.

Within the individual server view, navigate to the **Launch settings** tab.

The **Launch settings** tab is divided into two sections:

- General launch settings
- EC2 launch template

General launch settings

The **General launch settings** section allows you to control a variety of server-specific settings. Choose **Edit** to change the general settings.

Make your changes and then choose **Save settings** to finalize your changes.

Topics

- [Instance type right-sizing](#)
- [Start instance upon launching](#)
- [Copy Private IP](#)
- [Operating system licensing](#)
- [Transfer server tags](#)
- [Boot mode](#)

Instance type right-sizing

AWS Transform MGN launches a new instance type after every change of configuration on the source server, for example, added/removed disks, and added/removed RAM. When you use the instance type right-sizing feature MGN launches a test or cutover instance type that best matches

the OS, CPU, and RAM of your source server. Set this feature to **On** to enable or to **None** to disable, in which case MGN launches the instance type as configured in your Amazon EC2 launch template.

Enable instance type right-sizing if you want to determine the instance type that is launched in AWS for all your test or cutover servers.

Important considerations:

- The AWS instance type selected by AWS Transform MGN when this feature is activated overwrites the instance type defined in your EC2 launch template.
- Hardware changes and the resulting AWS instance type change may take up to 90 minutes to be processed by AWS Transform MGN.
- The T family instance type is not supported for right-sizing. If you want to use the T family, avoid using right-sizing.
- The available capacity for each Amazon EC2 instance type varies by Availability Zone and Region, and may be subject to your specific AWS account limits. For mission-critical workloads consider using [Reserved Instances](#) to guarantee capacity for specific instance types. Note that additional costs apply when reserving capacity.
- The right-sizing instance type selected by AWS Transform MGN appears on the **Server details** tab.

Supported instance families:

The service provides recommendations from these instance families, which are available in most target Regions:

- c5
- m5
- c4
- m4
- r5
- r4
- i3
- d2

Supported instance families for the Thailand and Malaysia Regions:

The *Asia Pacific (Thailand)* , and *Asia Pacific (Malaysia)* Regions support only these instance families:

- C6i
- m6i
- r6i

Start instance upon launching

Choose whether you want to start your test and cutover instances automatically upon launch or whether you want to launch them in a stopped state.

If you choose **Yes**, the test or cutover AWS instances are launched and started automatically upon test or cutover launch.

If you choose **No**, the instances are launched in a stopped state and you have to start the test or cutover AWS instance manually from the Amazon EC2 console.

Copy Private IP

Choose whether you want AWS Transform MGN to ensure that the private IP used by the test or cutover instance matches the private IP used by the source server.

AWS Transform MGN monitors the source server on an hourly basis to identify the private IP. MGN uses the private IP of the primary network interface.

The **No** option is chosen by default. Choose **No** if you do not want the private IP of the test or cutover instance to match that of the source machine.

Choose **Yes** if you want to use a private IP. The IP is shown in brackets next to the option.

Note

- Private IP is not supported for IPv6.
- Removing a private IP from a specific server's settings does not remove it from the launch template.

- If you chose **Yes**, ensure that the IP range of the subnet you set in the EC2 launch template includes the private IP address.
- If both the source server and the test or cutover instance share the same subnet through a VPN, then the source private IP is already in use, and the **Copy private IP** option should not be used.

Operating system licensing

Choose whether you want to Bring Your Own Licenses (BYOL) from the source server into the test or cutover instance.

Choose the **BYOL** option if you are migrating a Linux server. All Linux licenses are BYOL by default. Any RHEL, SUSE or Debian licenses are transferred in their current form to the migrated instance. Make sure that the terms of your licenses allow this license transfer.

Choose the **BYOL** option if you want to BYOL your Windows licenses. This sets up a dedicated host. All the licenses from the source Windows server are automatically transferred to the Test or Cutover instance. [Learn more about dedicated hosts.](#)

Important

If you activate BYOL licensing for Windows, you have to change the **Placement.tenancy** type in the EC2 launch template to **Host**. Otherwise, instance launch fails.

Note

- Windows Desktop Editions require BYOL – [note the specific restrictions for AWS Provided Licenses.](#)
- If you are using Windows Server Datacenter: Azure Edition, [note the specified restrictions for BYOL.](#)

Transfer server tags

Choose whether you want AWS Transform MGN to transfer any user-configured custom tags from your source servers onto your test or cutover instance.

If you choose **Yes**, server tags are transferred. These tags are attached to all source servers, all launched test and cutover instances, and all of the ephemeral resources that are created on your AWS Account during the normal operation of AWS Transform MGN. These resources include:

- EC2 instances
- Conversion groups
- Security groups
- Storage volumes (Amazon EBS or FSx for ONTAP)
- Snapshots

 Note

AWS Transform MGN automatically adds system tags to all resources.

 Note

Transfer server tags only copies tags associated with the source servers in the AWS Transform MGN console, and does not copy the EC2 source server tags (in case of AWS to AWS migration)

If you choose the **No** option, server tags are not transferred. You can always add tags from the Amazon EC2 console as described in [this EC2 article](#).

 Note

Tags that are added on the EC2 launch template take precedence over tags that are transferred directly from the source server.

Boot mode

Boot mode is automatically discovered from the source server. The target instance will launch using the same boot mode as the source. Changing this setting may cause the target instance to fail to boot.

UEFI limitations:

- UEFI boot is only available for Nitro instances.
- You must choose UEFI for any BYOL source server that is UEFI.
- UEFI is not supported on CentOS 6 and RHEL 6.

EC2 launch template

AWS Transform MGN uses EC2 launch templates to launch test and cutover EC2 instances for each source server.

The EC2 launch template is created automatically for each source server that is added to MGN upon the installation of the AWS Replication Agent.

Note

- MGN selects defaults to provide the best performance while migrating your servers to AWS. We recommend you review the EC2 launch template to ensure the selected templates are suitable for your use case.
- You cannot use the same template for multiple servers.
- The Launch template can only be edited from the Amazon EC2 console.
- Many EC2 launch template settings can be changed, but some may not be used by the AWS MGN launch process and some may interfere with IT. [Learn more about individual launch template settings.](#)

Important

- You must set the EC2 launch template you want to use with MGN as the **default** launch template.
- The EC2 launch template does not automatically set a specific subnet. As such, EC2 attempts to launch in a subnet within the default VPC. If you have removed your default VPC, EC2 fails to launch any instance for which there is no valid subnet specified. Ensure that you specify a subnet if that is the case, or MGN instance launch fails.

The MGN EC2 launch template panel shows a summary of the key template values. To view all the values or to change any of them:

1. Choose **Modify**.
2. When the **About modifying EC2 launch templates** dialog appears, choose **Modify**.

This redirects you to **EC2 > Launch templates > Modify template** in a new tab, where you'll be able to make any necessary changes.

Learn more about EC2 launch template settings and configuration options in [this EC2 article](#).

Selecting the default template

MGN uses the version of the Launch template that is marked as default.

In order to select the default launch template, on the **Modify template (Create new version)** page, under the **Launch template name and version description** category, open the **Source template** menu and choose the EC2 launch template you want to use as the default template from the drop-down menu.

Every time you modify the Launch template, a new version of the launch template is created. You are notified that the Launch template has been modified and that a new version (version number) has been created. Make sure to take note of the version number and the **Launch template ID** so that you could easily identify your launch template and version.

Note

It's good practice to delete versions of the launch template that you no longer need.

To set the new version of your launch template as the default:

1. Navigate back to the main **EC2 > Launch templates** page.
2. Choose your launch template by selecting the toggle to the left of the **Launch template ID**.
3. Open the Actions menu and choose **Set default version**.
4. Select the **Template version** from the drop-down menu and then choose **Set as default version**.

The Amazon EC2 console confirms the version change.

Launch template cleanup and fixing

AWS Transform MGN runs a mechanism every hour to ensure that the settings selected are correct. This mechanism can fix issues such as an incorrect instance type, but it cannot fix other settings and augmentations. Ensure that you follow the instructions in the following sections and do not change or edit any fields that should not be changed.

If you encounter any issues with the launch template, you can negate all of your changes and fix all issues rapidly by choosing the original default launch template that was first automatically created by MGN upon Agent installation.

Launch template key considerations

There are several key considerations when configuring your EC2 launch template. Review these key considerations as well as the [full launch settings](#) before creating your launch template.

1. **Instance Type** – Ensure that you select an instance type that matches the hardware requirements of your source server. AWS Transform MGN always uses the instance type that is set on the Amazon EC2 launch template unless the **Instance right-sizing** feature is activated.

Note

- If you change your instance type and do not deactivate the instance right-sizing feature, then AWS Transform MGN uses the instance type determined by the **Instance right-sizing** feature and not the instance type you chose in the EC2 launch template. MGN verifies the instance type once per hour, as a result, if you did not deactivate the instance right-sizing feature, the first time instance launch may still use the instance type you set in the EC2 launch template, but any subsequent launches use the right-sizing instance.
- The available capacity for each Amazon EC2 instance type varies by Availability Zone and Region, and may be subject to your specific AWS account limits. For mission-critical workloads consider using [Reserved Instances](#) to guarantee capacity for specific instance types. Note that additional costs apply when reserving capacity.

2. **Subnet** – You can select an existing subnet or create a new subnet.

 Note

Customers that do not have a default VPC must modify the EC2 launch template and explicitly define the subnet in which to launch. Failure to do so results in errors when launching test or cutover instances.

3. **Private IP** – If you use the **Copy private IP** feature, then do not add your own IP to the EC2 launch template. Private IP is not supported for IPv6.
4. **Private IP and Subnet** – Each subnet contains a CIDR block of IP ranges. If you use the **Copy private IP** feature, then ensure that this IP is included in the CIDR block range. Otherwise, instance launch fails.
5. **Private IP and ENI** – Make sure that you deactivate the **Copy private IP** feature if you wish to define an ENI to use on the EC2 launch template.
6. **Network interfaces** – The EC2 launch template only supports two network interfaces. If you require more than two network interfaces, you need to define them after the test or cutover instance has been launched. This can be done through a post launch action.

If you wish to use an Elastic IP, you must create an ENI to specify the IP and then edit the Network interfaces to use the ENI. Learn more about working with Elastic network interfaces in [the Amazon EC2 User Guide](#).

7. **Networking platform** – AWS Transform MGN only supports **Virtual Private Cloud (VPC)**. EC2-Classic is **not** supported. Do **not** add any security groups under the network platform.
8. **Custom device name** – Do not alter this field. AWS Transform MGN uses the device name as defined on the source server to map disks on the test or cutover instance. You can use this field to identify your disks.
9. **Disks** – You cannot add disks to the EC2 launch template. Any disks that are added that do not exist on the source machine are ignored by AWS Transform MGN.
- 10 **Launch template name** – Do not alter this field. AWS Transform MGN automatically names this field.
- 11 **System tag** – Do not alter this field. MGN automatically adds system tags that match the EC2 launch template to the specific source server. You can recognize which source server the launch template is matched with by the **ID** field.
- 12 **Automatic cleanup** – MGN deletes the EC2 launch template and launch configuration for machines that have been disconnected from AWS Transform MGN or machines for which the

cutover has been finalized 90 minutes after disconnect or cutover finalization. This aids in ensuring that your account does not surpass the AWS 5000 EC2 launch template limit.

 Note

When using FSx for ONTAP as the target storage type, the FSx for ONTAP split clone operation may take longer than 90 minutes. In this case, the EC2 instance and launch template are cleaned up within the standard 90-minute window, but the replica FSx for ONTAP volume may remain active until the split clone operation completes.

13.Volumes – For each storage volume, the service uses the user-selected values. If no matching volume exists in the launch template, the service uses the default value. If the launch template includes a volume that does not exist in the source server, the system disregards the specific volume.

If you delete the EC2 launch template, the service creates a new one with default values.

 Note

You can set a per-volume KMS key directly in the EC2 launch template's storage section. If set, the launch template KMS key takes precedence over the replication snapshot's key. Alternatively, you can configure encryption through the [EBS Encryption](#) section of the replication settings.

Full launch template setting review

This section reviews the entire EC2 launch template and identifies which fields should and should not be changed in order for the EC2 launch template to work with Application Migration Service. Editing or changing any fields that are marked as "do not edit" or "do not change" can cause AWS Transform MGN to not function.

- **Launch template name** – This name is automatically generated when the template is first created upon Agent installation. The name cannot be changed.
- **Template version description** – You can give the template any description you wish.

- **AMI** – Customers do not typically choose a specific AMI to include in the launch template. If you edit the launch template to use an existing AMI, the contents of the AMI are not used by AWS Transform MGN. If the AMI is not configured properly (licensing, flags, and more), then this may prevent the test or cutover instance launched from booting correctly or from being properly licensed.
- **Instance type** – You can select any instance type you want. The launch template shows the instance type suggested by AWS Transform MGN.
- **Key pair (login)** – **Do not** alter this field. Do not include a key pair with the launch template.
- **Networking platform** – Be sure to select **Virtual Private Cloud (VPC)**. **EC2-Classic** is not supported.
- **Security groups** – **Do not** add Security group here. This field should remain blank. You can add security groups later under **Network interface**.
- **Storage (volumes)** – This section shows all of the disks that you chose to replicate from your source server upon AWS Replication Agent installation.

 Important

Initial settings for EBS volumes are not derived from activity on the Source Server. Default values are chosen to give maximum performance on first launch.

Each disk is composed of the following fields:

- **Storage type** – Shows the target storage type (Amazon EBS or FSx for ONTAP) configured in the [replication settings](#).
- **Device name** – **Do not** change or edit this field. The device name shown here corresponds to the disk name on the source server. This field allows you to identify which disk is which.
- **Snapshot** – **Do not** change or edit this field. Snapshots should not be included in the launch template.
- **Size** – **Do not** change or edit this field.

- **Volume type** – You can select any volume type you want to use. AWS Transform MGN automatically sets **General Purpose SSD (gp3)** as the default. You may want to change the volume type to reduce costs. Ensure that you read the caveats in the [EBS documentation](#).
- **IOPS** – Set the number of I/O operations per second that the volume can support. You can select any number as long as it matches the [EBS guidelines](#).

- Provisioned IOPS SSD (io1) : 50 IOPS per GiB of storage
- Provisioned IOPS SSD (io2) : 500 IOPS per GiB of storage
- General Purpose SSD (gp3) : 500 IOPS per GiB of storage

AWS Transform MGN automatically provisions the maximum IOPS possible for the volume, based on the above ratio. This is to minimize the impact of the [performance penalty](#) when working with EBS volumes created from snapshots.

- **Throughput** – Set the throughput in MiB/s for the volume. This setting applies only to gp3 volumes. Refer to the [EBS volume types](#) documentation for supported values.
- **Add volume** – **Do not** use this functionality. You cannot add volumes to the source server through the launch template.
- **Remove (volume)** – **Do not** use this functionality. You cannot remove volumes from the source server through the launch template. If you do, MGN automatically creates a volume using the default volume settings.

 Important

The following storage parameters: **Volume initialization rate**, **EBS card index**, **Delete on termination**, and **KMS key** can only be configured per individual source server. Bulk editing of these parameters is not currently supported. To edit these parameters through the EC2 launch template console, see [Selecting the default template](#).

Volume initialization rate – Controls how fast a volume created from a snapshot is initialized, at a provisioned rate of 100–300 MiB/s. Use this for latency-sensitive workloads where you need the volume fully initialized quickly after launch. This setting is not supported on Outposts, Local Zones, or Wavelength Zones. AWS charges apply for provisioned initialization rate. If not set, no provisioned initialization is used and the volume initializes at default speed.

- **EBS card index** – Assigns a volume to a specific EBS card on instance types with multiple EBS controllers. Use this to spread I/O across cards for higher aggregate throughput. This setting

only has an effect on multi-card instance types (for example, `i4i.metal`, `r5b.24xlarge`). If not set, AWS automatically distributes volumes across available EBS cards.

- **Delete on termination** – Specifies whether an Amazon EBS volume is automatically deleted when the attached Amazon EC2 instance is terminated. Set this to **No** if you want the volumes to persist after instance termination, for example to preserve data for rollback or to reattach the volume to another instance. Set this to **Yes** for data volumes when you want to avoid orphaned volumes that continue incurring storage costs. This setting takes effect only after cutover is finalized and the instance is under your ownership. If not explicitly configured, the default behavior is **Yes** for root volumes and **No** for additional volumes.
- **Encrypted** – Set this to **Yes** if you are specifying a KMS key for the volume. If you are not using a per-volume KMS key, encryption is controlled through the [EBS Encryption](#) section of the replication settings.
- **KMS key** – You can specify a customer-managed KMS key to use for encrypting target volumes. If set, this key takes precedence over the key used during replication. If not set, the replication snapshot's KMS key is used. Ensure that your MGN launch role has `kms:CreateGrant`, `kms:Decrypt`, and `kms:GenerateDataKey` permissions on the specified key.
- **Resource tags** – You can add up to 50 tags. These are transferred to your test and cutover instances. Note that these tags may interfere with other tags that have already been added to the source server. Launch template tags always take precedence over tags set in the MGN console or tags manually assigned to the server.
- **Network interfaces** – The network interface is created by default based on your replication template. The network interface section is composed of the following fields:
 - **Device index** – **Do not** change or edit this field. The value should always be "0".
 - **Network interface** – Use this option only if you want to use a pre-existing ENI (Elastic Network Interface). The Launch Template overwrites certain ENI settings. Use this if you want to add an Elastic IP. You have to attach the Elastic IP to the ENI.

 Note

When selecting a pre-existing ENI, you must change the **Auto-assign public IP** value to **Don't include in launch template** for a successful target launch.

- **Description** – Add an optional description for the network interface (if chosen).
- **Subnet** – Choose the subnet. This is the subnet within which the network interface is located and the test or cutover instance is launched. AWS Transform MGN selects the default VPC subnet by default (if one exists).
- **Auto-assign public IP** – Choose whether you want the public IP to be auto-assigned.
- **Primary IP** – Use this field if you wish to use a private IP. The private IP you set in the **Copy private IP** field in the MGN launch settings is copied to this field.
- **Secondary IP** – Define a secondary IP, if needed.
- **IPv6 IPs** – Define IPv6 IPs if needed.
- **Security groups** – Choose a security group. If no security group is chosen, then the default VPC security group is used by default.

 Note

If you have selected FSx for ONTAP as the target storage in Replication settings, ensure you select a security group that allows your launched instances to connect to the FSx for ONTAP file system over iSCSI (TCP port 3260). See [Step 1: Configure security groups](#) in the [FSx for ONTAP configuration](#).

- **Delete on termination** – We suggest choosing "Yes". Choosing "No" makes this network interface a permanent ENI.
- **Elastic Fabric Adapter** – **Do not** change or edit this field.
- **Network card index** – **Do not** change or edit this field.
- **Add network interface** – Note that the EC2 launch template only supports two network interfaces. If you require more than two network interfaces, you need to define them after the test or cutover instance has been launched. This can be done through a post-launch action.
- **Advanced details** – In this section, we focus on the fields you should **not** change or edit to allow AWS Transform MGN to function properly. **Do not** change or edit any of the following fields:
 - RAM disk ID
 - Kernel
 - Nitro Enclave
 - Metadata accessible

Saving your EC2 launch template

Once you have finished editing your template, save it by choosing **Create template version** at the bottom of the template.

Launching test instances

After you have added all of your source servers and configured their launch settings, you are ready to launch a test instance. It is crucial to test the migration of your source servers to AWS before initiating a cutover to verify that your source servers function properly within the AWS environment.

Important

- It is a best practice to perform a test at least two weeks before you plan to migrate your source servers. This time frame allows you to identify potential problems and solve them, before the actual cutover takes place. After launching test instances, use either SSH (Linux) or RDP (Windows) to connect to your instance and ensure that everything is working correctly.
- When launching a test or cutover instance, you can launch up to 100 source servers in a single operation. Additional source servers can be launched in subsequent operations.

You can test one source server at a time, or simultaneously test multiple source servers. For each source server, you are informed of the success or failure of the test. You can test your source server as many times as you want. Each new test first deletes any previously launched Test instance and dependent resources. Then, a new test instance is launched, which reflects the most up-to-date state of the source server. After the test, data replication continues as before. The new and modified data on the source server is transferred to the Staging Area Subnet and not to the test instances that were launched during the test.

Note

- Windows source servers need to have at least 2 GB of free space to successfully launch a test instance.
- Take into consideration that once a test instance is launched resources are used in your AWS account and you will be billed for these resources. You can terminate the operation

of launched Test instances once you verify that they are working properly without impact to data replication.

Topics

- [Review ready for testing indicators](#)
- [Starting a test](#)
- [Reverting a test](#)
- [Marking as Ready for cutover](#)

Review ready for testing indicators

Before launching a Test instance, ensure that your source servers are ready for testing by looking for the following indicators on the **Source servers** page:

1. Under the **Migration lifecycle** column, the server should show **Ready for testing**.
2. Under the **Data replication status** column, the server should show the **Healthy** status.
3. Under the **Next step** column, the server should show **Launch test instance**

Starting a test

To launch a test instance for a single source server or multiple source servers:

1. Go to the **Source servers** page.
2. Check the box to the left of each server for which you want to launch a test instance.
3. Open the **Test and cutover** menu.
4. Under **Testing**, choose the **Launch test instances** option to launch a test instance for this server.
5. When the **Launch test instances for X servers** dialog appears, choose **Launch** to begin the test.

The AWS Transform MGN console indicates **Launch job started** when the test has started.

Choose **View job details** on the dialog to view the specific job for the test launch in the **Launch history** tab.

Successful test launch indicators

You can tell that the test instance launch started successfully through several indicators on the **Source servers** page.

1. The Alerts column shows the **Launched** status, indicating that a Test instance has been launched for this server.
2. The **Migration lifecycle** column shows **Test in progress**.
3. The **Next step** column shows **Complete testing and mark as 'Ready for cutover'**.

Reverting a test

After you have launched your test instances, open the Amazon EC2 console and SSH or RDP into your test instances to ensure that they function correctly. Validate connectivity and perform acceptance tests for your application.

If you encounter any issues and want to launch new Test instances, or if you are performing a scheduled test and plan to perform additional tests before cutover, you can revert the test. This reverts your source servers' **Migration lifecycle** status to **Ready for testing**, indicating that these servers still require additional testing before they are ready for cutover. During a revert, you also have the option to delete your Test instances for cost-saving purposes.

To revert a test:

1. Check the box to the left of every source server that has a launched Test instance for which you want to revert the test.
2. Open the **Test and cutover** menu.
3. Under **Testing**, choose **Revert to "ready for testing"**.
4. The **Revert testing for X servers** dialog appears. Select whether you want to terminate the launched instances used for testing. It is recommended to terminate these instances, as you will be charged for them even though you no longer need them. Check the **Yes, terminate launched instances (recommended)** box and choose **Revert**.

The AWS Transform MGN console indicates that testing has been reverted. The selected source servers' **Migration lifecycle** column shows the **Ready for testing** status, the **Next step** column shows **Launch test instance** and the launched Test instances are deleted if that option was selected.

Marking as Ready for cutover

If you are completely done with your testing and are ready for cutover, you can finalize the test. This changes your source servers' **Migration lifecycle** status to **Ready for cutover**, indicating that all testing is complete and that these servers are now ready for cutover. You also have the option to delete your Test instances for cost saving purposes.

To finalize a test:

1. Check the box to the left of every source server that has a launched Test instance for which you want to finalize the test.
2. Open the **Test and cutover** menu.
3. Under **Testing**, choose **Mark as "Ready for cutover"**.
4. When the **Mark X servers as "Ready for cutover"** dialog appears, select whether you want to terminate the launched instances used for testing. It is recommended to terminate these instances, as you will be charged for them even though you no longer need them. Check the **Yes, terminate launched instances (recommended)** box and choose **Continue**.

The AWS Transform MGN console confirms that the servers were marked as ready for cutover.

The AWS Transform MGN console indicates that testing has been finalized. The selected source servers' **Migration lifecycle** column shows the **Ready for cutover** status and the launched Test instances are deleted if that option was selected. The **Next step** column shows **Terminate launched instance; Launch cutover instance**.

You can now terminate the launched Test instance directly from the Amazon EC2 console as that instance is no longer needed (if you have not done so already through the MGN console). You can quickly access the Test instance by navigating to the specific servers > **Server details** > **Migration dashboard** > **Lifecycle** > **Launch status** and choosing **view in EC2 console**.

The Amazon EC2 console automatically searches for and displays the Test instance. Select the instance, open the **Instance state** menu, and choose **Terminate instance**.

Choose **Terminate**.

Launching cutover instances

Once you have finalized the testing of all of your source servers, you are ready for cutover. You should perform the cutover at a set date and time. The cutover migrates your source servers to the cutover instances on AWS.

Important

It is a best practice to perform a test at least two weeks before you plan to migrate your source servers. This time frame allows you to identify potential problems and solve them, before the actual migration takes place. After launching Test instances, use either SSH (Linux) or RDP (Windows) to connect to your instance and ensure that everything is working correctly.

You can cutover one source server at a time, or simultaneously cutover multiple source servers. For each source server, you are informed of the success or failure of the cutover. For each new cutover, AWS Transform MGN first deletes any previously launched Test instance and dependent resources. Then, it launches a new cutover instance which reflects the most up-to-date state of the source server. After the cutover, data replication continues as before. The new and modified data on the source server is transferred to the staging area subnet, and not to the cutover instances that were launched during the cutover.

Topics

- [Ready for cutover indicators](#)
- [Starting a cutover](#)
- [Reverting a cutover](#)
- [Finalizing a cutover](#)

Ready for cutover indicators

Before launching a cutover instance, ensure that your source servers are ready for cutover by looking for the following indicators on the **Source servers** page:

1. Under the **Migration lifecycle** column, the server should show **Ready for cutover** .
2. Under the **Data replication status** column, the server should show the **Healthy** status.

3. Under the **Next step** column, the server should show **Terminate launched instance; Launch cutover instance** if you have not terminated your latest launched test instance.
4. Alternatively, the Next step column shows **Launch cutover instance** if you have terminated your latest launched test instance.

Starting a cutover

To launch a cutover instance for a single source server or multiple source servers, go to the **Source servers** page and check the box to the left of each server you want to cutover.

Open the **Test and cutover** menu.

Under **Cutover**, choose the **Launch cutover instances** option.

The **Launch cutover instances for X servers** dialog appears. Choose **Launch** to begin the cutover.

In the **Source servers** page, the **Migration lifecycle** column shows **Cutover in progress** and the **Next step** column shows **Finalize cutover**.

The AWS Transform MGN console indicates **Launch job started** when the cutover has started.

Choose **View job details** on the dialog to view the specific job for the cutover launch in the **Launch history** tab.

Successful cutover launch indicators

You can tell that the cutover instance launch was started successfully through several indicators on the **Source servers** page.

1. The **Alerts** column displays **Launched**.
2. The **Migration lifecycle** column displays **Cutover in progress**.
3. The **Data replication status** displays **Healthy**.
4. The **Next step** column displays **Finalize cutover**.

Reverting a cutover

Once you have launched your cutover instances, open the Amazon EC2 console and SSH or RDP into your cutover instances to ensure that they function correctly. Validate connectivity and perform acceptance tests for your application.

Note

You should turn on Termination Protection after you have completed your testing and before you are ready to finalize the cutover. Learn more about enabling termination protection in [this Amazon EC2 article](#).

If you encounter any issues and want to launch new cutover instances, you can revert the cutover. This reverts your source servers' **Migration lifecycle** status to **Ready for cutover**, indicating that these servers have not undergone cutover. During a revert, you also have the option to delete your Cutover instances for cost-saving purposes.

To revert a cutover take the following steps:

1. Check the box to the left of every source server that has a launched cutover instance you want to revert.
2. Open the **Test and cutover** menu.
3. Under **Cutover**, choose **Revert to "ready for cutover"**.
4. This reverts your source servers' **Migration lifecycle** status to **Ready for cutover**, indicating that these servers have not undergone cutover.

When the **Revert cutover for X servers** dialog appears, choose **Revert**.

Finalizing a cutover

If you are completely done with your migration and performed a successful cutover, you can finalize the cutover. This changes your source servers' **Migration lifecycle** status to **Cutover complete**, indicating that the cutover is complete and that the migration has been performed successfully. In addition, this stops data replication and causes all replicated data to be discarded. All AWS resources used for data replication are terminated.

To finalize a cutover:

1. Check the box to the left of every source server that has a launched cutover instance you want to finalize.
2. Open the **Test and cutover** menu.
3. Under **Cutover**, choose **Finalize cutover**.
4. The **Finalize cutover for X servers** dialog appears. Choose **Finalize**. This changes your source servers' **Migration lifecycle** status to **Cutover complete**, indicating that the cutover is complete and that the migration has been performed successfully. In addition, this stops data replication and causes all replicated data to be discarded. All AWS resources used for data replication are terminated.

The AWS Transform MGN console indicates **Cutover finalized** when the cutover has completed successfully.

The AWS Transform MGN console automatically stops data replication for the source servers that were cutover to save resource costs. The selected source servers' **Migration lifecycle** column shows the **Cutover complete** status, the **Data replication** status column shows **Disconnected**, and the **Next step** column shows **Mark as archived**. The source servers have now been successfully migrated into AWS.

5. You can now archive your source servers that have launched cutover instances. Archiving removes these source servers from the main **Source servers** page, allowing you to focus on source servers that have not yet been cutover. You are still able to access the archived servers through filtering options.

To archive your cutover source servers:

- a. Check the box to the left of each source server for which the **Migration lifecycle** column states **Cutover complete**.
- b. Open the **Actions** menu and choose **Mark as archived**.
- c. When the **Archive X server** dialog appears, choose **Archive**.
- d. To see your archived servers, open the **Preferences** menu by choosing the gear button.

Toggle the **Show only archived servers** option and choose **Confirm**.

You are now able to see all of your archived servers. Untoggle the **Show only archived servers** option to show non-archived servers.

Review launch history

The **Launch history** tab allows you to track and manage all of the operations performed in AWS Transform MGN.

You can access the Launch History by choosing **Launch history** on the left-hand navigation menu.

Overview

The Launch History tab shows all of the operations (referred to as "Jobs") performed on your account. Each Job corresponds to a single operation (for example, Launch cutover instance, Launch test instance, etc.) Each Job is composed of one or more servers. The main Launch History view allows you to easily identify all key Job parameters, including:

- **Job ID** – The unique ID of the Job.
- **Job type** – The type of Job (Launch or Terminate)
- **Initiated by** – The command or action that initiated the job (for example, Launch cutover instances or Terminate launched instances)
- **Status** – The status of the Job (Pending, Completed, or Started)
- **Servers** – The number of servers that are included in the Job.
- **Start time** – The time the job was started.
- **Completed time** – The time the Job was completed (blank if the job was not completed).

You can sort the launch history by any column by choosing the column header. (for example, sorting by **Job ID**).

You can search for specific Jobs by any of the available fields within the **Find launch history by property or value** search bar.

Job details

To view a detailed breakdown of each individual job, choose the **Job ID** of the specific job.

The **Job details** view is composed of 3 sections:

Topics

- [Details](#)

- [Job log](#)
- [Jobs – Source servers](#)

Details

The **Details** section shows the same information as the main Job log page, including the **Type**, **Status**, **Initiated by**, **Start time**, and **Completed time**.

Job log

The Job log section shows a detailed log of all of the operations performed during the job.

Use this section to troubleshoot any potential issues and determine in which step of the launch process they occurred.

Use the **Filter job log by property or value** search bar to filter the job log.

You can filter by a variety of properties, including **Time**, **Event**, **Source server Id**, **Source server hostname**, **Conversion server instance ID**, **Test/cutover instance ID**, and **Error**.

You can filter by multiple values at once.

Jobs – Source servers

The **Source servers** section shows a list of all source servers involved in the job and their status.

You can use the **Filter source servers by property or value** search bar to filter by **Source server name** or **Status**.

Choose the **Source server name** of any source server from the list to open the Server Details view for that server. [Learn more about server details.](#)

Target storage configuration

AWS Transform MGN supports two target storage types for replication. You can select the storage type that best suits your workload requirements:

- **Amazon Elastic Block Store (Amazon EBS)** – The default storage type. For configuration details, see [Amazon EBS configuration](#).
- **Amazon FSx for NetApp ONTAP** – Enterprise file storage with NetApp ONTAP capabilities. For setup instructions, see [FSx for ONTAP configuration](#).

Amazon EBS configuration

Amazon Elastic Block Store (Amazon EBS) is the default target storage type for AWS Transform MGN. This guide covers how to configure Amazon EBS volume types, encryption settings, and per-server disk type customization.

Replication template settings

These settings are configured in the replication template and apply to all source servers by default.

Amazon EBS volume type

Choose the default **Amazon EBS volume type** to be used by the replication servers for large disks.

Each disk has minimum and maximum sizes and varying performance metrics and pricing. Learn more about Amazon EBS volume types in [this Amazon EBS article](#).

The best practice is to not change the default Amazon EBS volume type, unless there is a business need for doing so.

Note

This option only affects disks over 500 GiB (by default, smaller disks always use Magnetic HDD volumes).

The default **Lower cost, Throughput Optimized HDD (st1)** option uses slower, less expensive disks.

You might want to use this option if:

- You want to keep costs low
- Your large disks do not change frequently
- You are not concerned with how long the initial sync process will take

The **Faster, General Purpose SSD (gp3)** option uses faster, but more expensive disks.

You might want to use this option if:

- Your source server has disks with a high write rate or if you want faster performance in general
- You want to speed up the initial sync process
- You are willing to pay more for speed

Note

You can customize the Amazon EBS volume type used by each disk within each source server in that source server's settings. [Learn more about changing individual source server volume types.](#)

Amazon EBS encryption

Choose whether to use the default or custom Amazon **EBS encryption**. This option will encrypt your replicated data at rest on the Staging Area Subnet disks and the replicated disks.

- **Default** – The default Amazon EBS encryption Volume Encryption Key will be used (which can be an Amazon EBS-managed key or a CMK).
- **Custom** – You will need to enter a custom customer-managed key (CMK) in the regular key ID format.

If you select the **Custom** option, the **EBS encryption key** box appears. Enter the ARN or key ID of a customer-managed CMK from your account or another AWS account. Enter the encryption key (such as a cross-account KMS key) in the regular key ID format (KMS key example: 123abcd-12ab-34cd-56ef-1234567890ab).

To create a new AWS Key Management Service key, choose **Create an AWS KMS key**. You will be redirected to the Key Management Service (KMS) Console where you can create a new key to use.

Learn more about Amazon EBS Volume Encryption in [this Amazon EBS article](#).

Important

Reversing the encryption option after data replication has started will cause data replication to start from the beginning.

Using an AWS KMS Customer Managed Key (CMK) for encryption

If you decide to use a Customer Managed Key (CMK), or if your default Amazon EBS encryption key is a CMK, you will need to add additional permissions to the key to allow AWS Transform MGN to use it.

To modify the existing key policy using the AWS Management Console *policy view*.

1. Navigate to the AWS KMS Console and select the AWS KMS key you plan to use with AWS MGN.
2. Scroll to **Key policy** and choose **Switch to policy view**.
3. Choose **Edit** and add the following JSON statements to the **Statement** field.

```
{
  "Sid": "Allow AWS Services permission to describe a customer managed key for
encryption purposes",
  "Effect": "Allow",
  "Principal": {
    "AWS": "arn:aws:iam::$ACCOUNT_ID:root"
  },
  "Action": [
    "kms:DescribeKey"
  ],
  "Resource": "*",
  "Condition": {
    "StringEquals": {
      "kms:CallerAccount": [
        "$ACCOUNT_ID"
      ]
    }
  },
  "Bool": {
    "aws:ViaAWSService": "true"
  }
}
```

```

    },
    {
      "Sid": "Allow AWS MGN permissions to use a customer managed key for EBS
encryption",
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::$ACCOUNT_ID:root"
      },
      "Action": [
        "kms:CreateGrant"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "kms:CallerAccount": [
            "$ACCOUNT_ID"
          ],
          "kms:GranteePrincipal": [
            "arn:aws:iam::$ACCOUNT_ID:role/aws-service-role/mgn.amazonaws.com/
AWSServiceRoleForApplicationMigrationService"
          ]
        },
        "ForAllValues:StringEquals": {
          "kms:GrantOperations": [
            "CreateGrant",
            "DescribeKey",
            "Encrypt",
            "Decrypt",
            "GenerateDataKey",
            "GenerateDataKeyWithoutPlaintext"
          ]
        },
        "Bool": {
          "aws:ViaAWSService": "true"
        }
      }
    },
    {
      "Sid": "Allow EC2 to use this key on behalf of the current AWS Transform MGN
user, during target launches",
      "Effect": "Allow",
      "Principal": {
        "AWS": [
          "arn:aws:iam::$ACCOUNT_ID:root",

```



```

    "arn:aws:iam::${ACCOUNT_ID}:role/aws-service-role/mgn.amazonaws.com/
    AWSServiceRoleForApplicationMigrationService"
  ],
  "Action": [
    "kms:ReEncrypt*",
    "kms:GenerateDataKey*"
  ],
  "Resource": "*",
  "Condition": {
    "StringEquals": {
      "kms:CallerAccount": [
        "${ACCOUNT_ID}"
      ],
      "kms:ViaService": "ec2.${REGION}.amazonaws.com"
    }
  }
}

```

Important

- Replace `$ACCOUNT_ID` with the AWS account ID you are migrating into.
- Replace `$REGION` with the AWS Region you are migrating into.
- The last statement can be made stricter by ensuring the principal refers to users who are going to perform [StartTest](#) or [StartCutover](#) API calls

4. Choose **Save changes**.

Note

If you are using a Customer Managed Key (CMK) from another account, you need to take an additional step from within that account to allow the service to use the CMK.

From the account in which you want to stage MGN replication servers, create a grant that delegates the relevant permissions to the appropriate service-linked role. The Grantee Principal element of the grant is the ARN of the appropriate service-linked role. The key-id is the ARN of the key.

The following is an example [create-grant](#) CLI command that gives the service-linked role named **AWSServiceRoleForApplicationMigrationService** in account 111122223333 permissions to use the customer-managed key in account 444455556666.

```
aws kms create-grant \  
--region us-west-2 \  
--key-id arn:aws:kms:us-  
west-2:444455556666:key/1a2b3c4d-5e6f-1a2b-3c4d-5e6f1a2b3c4d \  
--grantee-principal arn:aws:iam::111122223333:role/aws-service-role/  
mgn.amazonaws.com/AWSServiceRoleForApplicationMigrationService \  
--operations "Encrypt" "Decrypt" "ReEncryptFrom" "ReEncryptTo" "GenerateDataKey"  
"GenerateDataKeyWithoutPlaintext" "DescribeKey" "CreateGrant"
```

For this command to succeed, the user making the request must have permissions for the **CreateGrant** action.

Server-level settings

These settings can be customized per source server or group of source servers.

Change staging disk type

You can change the Amazon EBS volume disk type for each disk or for a group of disks.

To change the Amazon EBS volume disk type, select the circle to the left of each disk name and choose **Change staging disk type**.

On the **Change staging disk type** dialog, select the type of Amazon EBS volume to use for the disk or group of disks.

Select the **AUTO** option if you want AWS Transform MGN to automatically select the most cost-effective Amazon EBS volume disk type for each disk based on the disk size and type based on the option you defined in the [Replication settings](#) (either the default **Lower cost, Throughput Optimized HDD (st1)** option or the **Faster, General Purpose SSD (gp3)** option).

AWS Transform MGN uses a single replication server per 15 source disks. Selecting the **AUTO** option ensures that the fewest number of replication servers are used, resulting in increased cost savings.

Note

AWS Transform MGN always uses Amazon EBS magnetic volumes for disks that are under 500 GiB in size when the **AUTO** option is selected.

If you do not want AWS Transform MGN to automatically select a disk, you can select a disk manually. Select the disk type from the **EBS volume type** menu.

Note

When replicating into an AZ, ensure that the AZ supports the staging disk type chosen.

For certain disks, you can configure the amount of IOPS to be allocated per GB of disk space under **IOPS**. You can allocate up to 50 IOPS per GB. 64,000 IOPS are available for Nitro-based instances. Other instances are guaranteed up to 32,000 IOPS. The maximum IOPS per instance is 80,000.

Choose **Change** to confirm the change.

For **General Purpose SSD (gp3)** disks, you can also set the **Throughput**. General Purpose SSD (gp3) volumes have a baseline performance of 125 MiB/s. You can provision additional throughput of 0.25 MiB/s per provisioned IOPS up to a maximum of 1,000 MiB/s (at 4,000 IOPS or higher).

Choose **Change** to confirm the change.

FSx for ONTAP configuration

Overview

This page provides step-by-step instructions for configuring Amazon FSx for NetApp ONTAP (FSx for ONTAP) as a storage migration target for AWS Transform MGN (MGN) when migrating to AWS. With this setup, you can use the enterprise file storage capabilities of FSx for ONTAP for your migrated workloads. This page assumes that you are familiar with FSx for ONTAP. For detailed FSx for ONTAP setup instructions, see the [FSx for ONTAP Getting Started Guide](#).

FSx for ONTAP as a target storage type is available in all AWS Regions where both MGN and FSx for ONTAP are available. This storage type is not available in Local Zones. For supported regions, see [MGN supported regions](#) and [FSx for ONTAP availability by Region](#).

⚠ Important

Do not rename or modify MGN-managed FSx for ONTAP resources (LUNs, igroups, snapshots). Changes to these resources disrupt migration and require restarting the migration from the beginning.

Known limitations

- **FSx for ONTAP backups can block volume cleanup.** FSx for ONTAP file systems have automatic backups enabled by default. Backups taken on target volumes can prevent MGN from deleting replication volumes on Finalize cutover/Disconnect from service migration stage. See [Troubleshooting replication volume not deleted after Finalize cutover](#).
- **Multiple LUNs per volume.** MGN creates one volume per source server on the FSx for ONTAP file system and places each disk as a separate LUN within that volume. For example, a source server with 3 disks results in one volume with 3 LUNs. The ONTAP best practice is a 1:1 relationship (one volume per LUN), which allows per-volume features such as snapshots, tiering policies, and storage efficiency to be configured independently per disk. As a workaround, you can use the ONTAP [lun move start](#) command to relocate LUNs into dedicated volumes after migration. This operation is non-disruptive and does not require iSCSI reconfiguration on the host.
- **Agent-based replication only.** MGN supports FSx for ONTAP as a target storage type only with agent-based replication.
- **Up to 5 file systems per account.** MGN supports migrating data into up to 5 FSx for ONTAP file systems concurrently per account. If you have more file systems, migrate in phases. For more information about FSx for ONTAP quotas, see [FSx for ONTAP quotas](#). For MGN service quotas, see [MGN endpoints and quotas](#).
- **igroup limits per file system.** MGN creates one igroup per source server during replication and one per target instance at launch. FSx for ONTAP supports up to 256 igroups (Single-AZ) or 512 igroups (Multi-AZ). Plan the number of source servers per file system accordingly.
- **ONTAP configurations not migrated.** If you are migrating from an existing ONTAP storage system, source ONTAP configurations (such as access permissions, quotas, snapshot policies, and schedules) are not migrated automatically. You must reconfigure these settings on the target FSx for ONTAP file system after migration.
- **No mixed storage per server.** All data volumes from a source server use the same storage type (either Amazon EBS or FSx for ONTAP). You cannot mix storage types for different disks on the same server. The boot volume is always stored on Amazon EBS.

Prerequisites

Before integrating FSx for ONTAP with MGN, ensure the following:

- **MGN Setup:** MGN initialized in your AWS account with agent-based replication.

Important

If you initialized MGN before FSx for ONTAP support was available, you must reinitialize the service to create the required AWS managed roles. In the MGN console, navigate to **Settings** → **Replication template** and choose **Reinitialize Service Permissions**. For details on these roles and their managed policies, see [AWS Transform MGN managed policies](#).

- **VPC Configuration:** FSx for ONTAP and MGN instances must be in the same AWS account and Region. They can use the same VPC or different VPCs in which both the source and target are routable. IPv4 connectivity is required. As a best practice, use a Multi-AZ FSx for ONTAP file system when your applications require resiliency across Availability Zones. Deploy your target EC2 instances in the same AZ as the file system's preferred file server to minimize latency. If your applications do not require cross-AZ resiliency, you can use a Single-AZ FSx for ONTAP file system. Deploy your target EC2 instances in the same AZ as the file system to minimize latencies and avoid cross-AZ data transfer charges.
- **OS package repository access:** MGN automatically installs iSCSI initiator and multipath packages on replication servers and launched instances. Ensure that both the staging area subnet and the launch subnet have outbound access to OS package repositories (for example, through a NAT gateway or internet gateway). For the full list of required URLs, see [Network requirements for FSx for ONTAP](#).

Step 1: Configure security groups

To enable MGN to work with FSx for ONTAP, you must create two security groups that cross-reference each other:

- **MGN-Instances-SG.** Attached to the EC2 instances that MGN launches (test and cutover).
- **FSx-ONTAP-SG.** Attached to the FSx for ONTAP file system. Controls inbound traffic from MGN-launched instances.

Because FSx-ONTAP-SG references MGN-Instances-SG as the source in its inbound rules, only MGN-launched instances can reach the file system. All other traffic is denied by default.

1.1 MGN instances security group

Create this security group in the VPC where MGN will launch target instances.

Important

If you use different VPCs for replication and launch, create two security groups with distinct names for clarity. For example, use MGN-Replication-SG (in the staging VPC) and MGN-Launch-SG (in the launch VPC).

Steps to create

1. Navigate to the **Amazon VPC Console** → **Security Groups** → **Create security group**.
2. Configure the following settings:
 - Security group name: MGN-Instances-SG
 - Description: Security group for instances launched by MGN to allow communication with FSx for ONTAP
 - VPC: Choose the target VPC where MGN will launch instances.
3. **Inbound Rules:** The only required inbound rule is port 1500 for MGN data replication from source servers. You can optionally add rules for administrative access to your instances (for example, SSH on port 22 or RDP on port 3389 from your corporate network).

Type	Protocol	Port Range	Source	Description
Custom TCP	TCP	1500	<i>Source server CIDR</i>	Data replication from source servers

4. **Outbound Rules:** The default outbound rule (All traffic → 0.0.0.0/0) is sufficient. If you restrict outbound rules, add at minimum the following rules. Reference the FSx for ONTAP security group (created in the next step) as the destination:

Type	Protocol	Port Range	Destination	Description
iSCSI	TCP	3260	FSx-ONTAP-SG	iSCSI access to FSx for ONTAP
HTTPS	TCP	443	FSx-ONTAP-SG	ONTAP REST API / Management

5. Choose **Create security group**.

1.2 FSx for ONTAP security group

You associate this security group with the FSx for ONTAP file system. Use this security group to control which resources can communicate with the file system and to ensure that MGN-launched instances have the necessary access.

Steps to create

1. Navigate to the **Amazon VPC Console** → **Security Groups** → **Create security group**.
2. Configure the following settings:
 - Security group name: FSx-ONTAP-SG
 - Description: Security group for FSx for ONTAP file system to allow inbound access from MGN-launched instances
 - VPC: Choose the target VPC used for the FSx for ONTAP file system.
3. **Inbound Rules:** Add the following rules. The table is organized into two groups:
 - **Migration traffic (iSCSI).** Required for MGN data replication and launch. Reference MGN-Instances-SG as the source.
 - **Management access (SSH, HTTPS).** Optional rules for ONTAP CLI and REST API access from MGN-launched instances (for example, for troubleshooting or manual configuration). Reference MGN-Instances-SG as the source.
 - **MGN service traffic (HTTPS).** Required for MGN to access the FSx for ONTAP REST API during replication and launch. Use the CIDR blocks of the preferred and standby subnets where the file system is deployed. You can find these CIDRs in the **FSx for ONTAP Console** under your file system's **Network & security** tab, or in the **VPC Console** → **Subnets** by looking up the subnet IDs. You can narrow this scope after the initial setup is complete.

Type	Protocol	Port Range	Source	Description
Migration traffic				
iSCSI	TCP	3260	MGN-Instances-SG	Allow iSCSI from MGN instances
Management access (optional)				
SSH	TCP	22	MGN-Instances-SG	ONTAP CLI management from MGN instances
HTTPS	TCP	443	MGN-Instances-SG	ONTAP REST API management from MGN instances
MGN service traffic				
HTTPS	TCP	443	<i>FSx preferred subnet CIDR</i>	MGN access to ONTAP REST API
HTTPS	TCP	443	<i>FSx standby subnet CIDR</i>	MGN access to ONTAP REST API

4. **Outbound Rules:** The default outbound rule (All traffic → 0.0.0.0/0) is sufficient.

5. Choose **Create security group**.

Step 2: Create FSx for ONTAP file system

MGN can work with an existing FSx for ONTAP file system or a new one. If you do not already have an FSx for ONTAP file system, create one in the same AWS account and Region where MGN will launch target instances.

High-level steps

1. Navigate to the Amazon FSx for NetApp ONTAP Console and create a new FSx for ONTAP file system.

2. Choose **Standard create** option.
3. Choose deployment type (Multi-AZ or Single-AZ).
4. Configure storage capacity and throughput based on your workload requirements.
5. Choose the VPC and subnets for FSx for ONTAP deployment.
6. Choose FSx-ONTAP-SG security group created in [the section called "1.2 FSx for ONTAP security group"](#).
7. Configure FSx for ONTAP admin account and password.
8. Configure a Storage Virtual Machine (SVM).
9. Wait for the file system to reach **Available** status (approximately 30-45 minutes).

Important for MGN integration

- **For Multi-AZ deployments:** you must specify an Endpoint IPv4 address range (not unallocated or floating) that is outside your VPC CIDR. Use the "Enter an IPv4 address range" option and provide a specific range within RFC 1918 private address space (for example, 192.168.1.0/24). This is required for MGN integration to ensure consistent endpoint addressing.
- **Storage capacity:** MGN uses FSx for ONTAP storage for replication, conversion, and cutover. These processes require temporary storage on the file system. Ensure that sufficient space is available on the FSx for ONTAP file system and increase capacity if needed. As a guideline, provision 3x the size of the planned migration data. The 3x factor accounts for three concurrent storage consumers during migration: the replicated data, the converted volumes used for launch, and the original volumes pending deletion. Volume deletion in FSx for ONTAP is a background operation, so freed capacity is not available immediately after deletion, so headroom must be provisioned upfront. As a best practice, keep the file system at or below 80% SSD capacity utilization throughout the migration. For second-generation file systems ([Single-AZ 2](#) and [Multi-AZ 2](#) deployment types), you can decrease storage capacity after migration is complete. For more information, see [Managing storage capacity and provisioned IOPS](#).
- **Throughput capacity:** Higher throughput capacity reduces migration time. Throughput is selected from supported values and is a billable dimension. To size throughput for migration, sum the average read throughput and write throughput across all source servers being migrated to the file system, add 15% headroom, and round up to the next supported value. Plan your throughput capacity before starting migration, as changes

take time to take effect. You can reduce throughput after migration is complete. For more information, see [Managing throughput capacity](#) and [FSx for ONTAP performance](#).

Disable Autonomous Ransomware Protection (ARP)

If ONTAP ARP is enabled on the file system, disable it before migration. ARP can prevent MGN from deleting replication volumes on Finalize cutover/Disconnect from service migration stage. For more information, see [Enabling Autonomous Ransomware Protection](#).

For detailed instructions on creating and configuring FSx for ONTAP file systems, see [Creating FSx for ONTAP file systems](#).

Step 3: Configure certificate-based authentication

Certificate-based authentication is required for MGN to access the ONTAP REST API. MGN handles TLS validation internally using AWS Certificate Authorities.

Note

MGN does not use CHAP for iSCSI. iSCSI access is controlled via security groups, and MGN authenticates to the ONTAP management API using client certificates as described in this section.

Create client certificate for API authentication

Generate a client certificate that FSx for ONTAP will require and MGN will use to authenticate to the ONTAP REST API. You have several options:

Certificate options

Option	Use Case	Documentation
Self-Signed Certificate	Testing/Development	Generating a self-signed certificate for FSx for ONTAP

Option	Use Case	Documentation
AWS Private Certificate Authority	Production (Recommended)	AWS Private Certificate Authority
External Certificate Authority	Production (Enterprise PKI)	Use your organization's CA process

Note

The private key must be in PKCS#8 format (-----BEGIN PRIVATE KEY-----). If your key starts with -----BEGIN RSA PRIVATE KEY----- (PKCS#1), convert it:

```
[~]$ openssl pkcs8 -topk8 -inform PEM -outform PEM -nocrypt \
  -in fsx-mgn-client.key -out fsx-mgn-client.key.pk8
[~]$ mv fsx-mgn-client.key.pk8 fsx-mgn-client.key
```

Install client certificate on FSx for ONTAP

Install the client certificate on the FSx for ONTAP to enable certificate-based authentication. In the following commands, replace *vserver_name* with your file system ID (e.g., FsxId08f0e724d292c729c). You can find this in the FSx for ONTAP console under your file system's details.

1. Connect to the file system's management endpoint. Log in to an EC2 instance in the same VPC as the FSx for ONTAP file system, then use the `fsxadmin` user to SSH into the file system's management endpoint IP address or DNS name:

```
[~]$ ssh fsxadmin@file-system-management-endpoint-ip-address
```

For more information, see [Managing file systems with the ONTAP CLI](#).

2. Install the client CA certificate:

```
FsxId0123456::> security certificate install -type client-ca \
  -vserver vserver_name -cert-name my-client-ca
```

```
# Paste the contents of ca.crt when prompted
# Press Enter when done
```

Verify the certificate was installed:

```
FsxId0123456::> security certificate show -vserver vserver_name -type client-ca
```

3. Create the user with certificate authentication. For more information, see [Creating ONTAP users](#).

```
FsxId0123456::> security login create -vserver vserver_name \
  -user-or-group-name cert_usr -application http \
  -authentication-method cert -role fsxadmin
```

Verify the login was created:

```
FsxId0123456::> security login show -vserver vserver_name \
  -user-or-group-name cert_usr
```

Test certificate-based authentication

Log in to an EC2 instance in the same VPC as the FSx for ONTAP file system, then run the following command to test certificate authentication:

Download the FSx for ONTAP certificate bundle (change the region in the URL):

```
[~]$ curl https://fsx-aws-certificates.s3.amazonaws.com/bundle-region.pem \
  -o bundle-region.pem
```

Test authentication using the certificate:

```
[~]$ curl -sS --cacert bundle-region.pem \
  --cert fsx-mgn-client.crt \
  --key fsx-mgn-client.key \
  https://management.fs-xxxxx.fsx.region.amazonaws.com/api/cluster
```

Expected: JSON response with cluster information. If you see 401 Unauthorized, verify certificate installation and login creation.

Step 4: Store certificates in AWS Secrets Manager

Store the client certificate and private key in AWS Secrets Manager (Secrets Manager). MGN will retrieve these credentials using the Secret ARN.

Required secret format:

MGN expects the secret to contain exactly two keys:

- `cert`: The client certificate content (`fsx-mgn-client.crt`)
- `key`: The private key content (`fsx-mgn-client.key`)

Store using AWS Console:

1. Navigate to **Secrets Manager** in the AWS Console.
2. Choose **Store a new secret**.
3. Choose **Other type of secret**.
4. Add key-value pairs with exact key names as key/value (not plain text):
 - `cert` – content of your `fsx-mgn-client.crt`
 - `key` – content of your `fsx-mgn-client.key`
5. Choose **Next**.
6. On the **Configure secret** page, under **Tags**, add a tag with key `AWSApplicationMigrationServiceManaged` and value `True`.
7. Choose **Next** → **Store**.
8. **Copy the Secret ARN**. You need this for MGN configuration.

Important

- Use `cert` (not certificate).
- Use `key` (not `private_key`).
- Do NOT include a `username` field.

Example Secret ARN:

```
arn:aws:secretsmanager:us-east-1:123456789012:secret:mgn/fsx/ontap-api-  
certificate-AbCdEf
```

Step 5: Configure MGN replication settings

Configure MGN to use the FSx for ONTAP REST API certificate stored in Secrets Manager.

Configure replication template by using the AWS Console

Important

Changing the storage provider for a source server that is already replicating terminates current replication and restarts the replication process from the beginning.

1. Navigate to **MGN** console.
2. Under **Settings**, choose **Replication template**.
3. Choose **Edit**.
4. Choose the required target subnet (subnet that can communicate with FSx for ONTAP and has outbound access to OS package repositories).
5. Choose **FSx for ONTAP configuration**.
6. Enter the following configuration:
 - Choose **AWS FSx for ONTAP** as a default storage type.
 - **Storage Virtual Machine (SVM) ID**: choose from the list.
 - **FSx Storage Secret ARN**: enter the Secret ARN you copied in [the section called “Step 4: Store certificates in AWS Secrets Manager”](#).
7. Choose the MGN-Instances-SG security group (created in [the section called “1.1 MGN instances security group”](#)) to allow iSCSI traffic to FSx for ONTAP.
8. Choose **Save changes**.

Note

Migration Acceleration Program (MAP) 2.0 tags are applied to the FSx for ONTAP file system but not to individual FSx for ONTAP volumes.

Step 6: Configure launch template and launch settings

The target instance must establish iSCSI connectivity to the FSx for ONTAP SVM over the network.

Requirements:

- Choose the required target subnet. This subnet must have network connectivity to the FSx for ONTAP file system (ports 3260 and 443) and outbound access to OS package repositories. For connectivity details, see [Network requirements for FSx for ONTAP](#).
- Modify the source server's launch template to include the MGN-Instances-SG security group (see [the section called "Step 1: Configure security groups"](#)).

Step 7: Enable volume integrity validation (recommended)

Enable the [Volume integrity validation](#) post-launch action to automatically verify iSCSI connectivity and multipath mount configuration after each test or cutover launch. For FSx for ONTAP migrations, this action validates that all expected iSCSI volumes are connected, mounted, and accessible through multipath.

To enable this action, see [Post-launch settings](#).

Post-migration optimization

After successful cutover, optimize your FSx for ONTAP deployment for ongoing operations.

Configure backup strategy

After migration, verify that your FSx for ONTAP backup strategy covers the migrated data. Review automatic backup settings and retention policies for your file system, and confirm that migrated volumes are included in your backup schedule. For more information, see [Working with backups](#).

Re-enable Autonomous Ransomware Protection (ARP)

If you disabled ONTAP ARP before migration, re-enable it after cutover is complete. For more information, see [Enabling Autonomous Ransomware Protection](#).

Generating a self-signed certificate for FSx for ONTAP

To generate a self-signed CA and client certificate, save the following script as `generate-fsx-mgn-certs.sh` and run it:

```
[~]$ chmod +x generate-fsx-mgn-certs.sh
./generate-fsx-mgn-certs.sh "YourOrg" "US"
```

```
#!/bin/bash
set -e
ORG="${1:-YourOrg}"
COUNTRY="${2:-US}"
DAYS_CA="${3:-3650}"
DAYS_CLIENT="${4:-365}"

# Generate CA private key and certificate
openssl genrsa -out ca.key 4096
openssl req -new -x509 -key ca.key -out ca.crt -days "$DAYS_CA" \
  -subj "/CN=FSx-ONTAP-Client-CA/O=$ORG/C=$COUNTRY" \
  -addext basicConstraints=critical,CA:TRUE \
  -addext keyUsage=critical,keyCertSign,cRLSign \
  -addext subjectKeyIdentifier=hash

# Generate client key
openssl genrsa -out fsx-mgn-client.key 2048

# Create OpenSSL config
cat > openssl-client.cnf << EOF
[ req ]
default_bits = 2048
prompt = no
default_md = sha256
distinguished_name = dn
req_extensions = req_ext
[ dn ]
CN = cert_usr
O = $ORG
C = $COUNTRY
[ req_ext ]
keyUsage = critical, digitalSignature, keyEncipherment
extendedKeyUsage = clientAuth
subjectKeyIdentifier = hash
[ usr_cert ]
keyUsage = critical, digitalSignature, keyEncipherment
extendedKeyUsage = clientAuth
authorityKeyIdentifier = keyid,issuer
subjectKeyIdentifier = hash
EOF
```



```
# Create CSR and sign it
openssl req -new -key fsx-mgn-client.key -out fsx-mgn-client.csr \
  -config openssl-client.cnf
openssl x509 -req -in fsx-mgn-client.csr \
  -CA ca.crt -CAkey ca.key -CAcreateserial \
  -out fsx-mgn-client.crt -days "$DAYS_CLIENT" \
  -extfile openssl-client.cnf -extensions usr_cert

# Convert key to PKCS#8 format
openssl pkcs8 -topk8 -inform PEM -outform PEM -nocrypt \
  -in fsx-mgn-client.key -out fsx-mgn-client.key.pk8
mv fsx-mgn-client.key.pk8 fsx-mgn-client.key

# Verify
openssl verify -CAfile ca.crt fsx-mgn-client.crt
```

Output files:

- `fsx-mgn-client.crt` – Client certificate
- `fsx-mgn-client.key` – Private key (PKCS#8 format)
- `ca.crt` – CA certificate (install on FSx for ONTAP)

Manual steps (alternative):

If you prefer to run each step individually:

1. Create a Certification Authority (CA):

```
# Generate CA private key
[~]$ openssl genrsa -out ca.key 4096

# Create self-signed CA certificate
[~]$ openssl req -new -x509 -key ca.key -out ca.crt -days 3650 \
  -subj "/CN=FSx-ONTAP-Client-CA/0=YourOrg/C=US" \
  -addext basicConstraints=critical,CA:TRUE \
  -addext keyUsage=critical,keyCertSign,cRLSign \
  -addext subjectKeyIdentifier=hash
```

2. Generate a client key:

```
[~]$ openssl genrsa -out fsx-mgn-client.key 2048
```

3. Create an openssl-client.cnf file:

```
[ req ]
default_bits      = 2048
prompt           = no
default_md       = sha256
distinguished_name = dn
req_extensions    = req_ext

[ dn ]
CN = cert_usr
O  = YourOrg
C  = US

[ req_ext ]
keyUsage          = critical, digitalSignature, keyEncipherment
extendedKeyUsage  = clientAuth
subjectKeyIdentifier = hash

[ usr_cert ]
keyUsage          = critical, digitalSignature, keyEncipherment
extendedKeyUsage  = clientAuth
authorityKeyIdentifier = keyid,issuer
subjectKeyIdentifier = hash
```

4. Create a new certificate signing request (CSR):

```
[~]$ openssl req -new -key fsx-mgn-client.key -out fsx-mgn-client.csr \
    -config openssl-client.cnf
```

5. Sign the CSR with your CA:

```
[~]$ openssl x509 -req -in fsx-mgn-client.csr \
    -CA ca.crt -CAkey ca.key -CAcreateserial \
    -out fsx-mgn-client.crt -days 365 \
    -extfile openssl-client.cnf -extensions usr_cert
```

6. Convert key to PKCS#8 format:

```
[~]$ openssl pkcs8 -topk8 -inform PEM -outform PEM -nocrypt \
```

```
-in fsx-mgn-client.key -out fsx-mgn-client.key.pk8  
[~]$ mv fsx-mgn-client.key.pk8 fsx-mgn-client.key
```

7. Verify the certificate:

```
[~]$ openssl verify -CAfile ca.crt fsx-mgn-client.crt
```

Monitoring Application Migration Service

Monitoring is an important part of maintaining the reliability, availability, and performance of AWS Transform MGN and your other AWS solutions. AWS provides the following monitoring tools to watch Application Migration Service, report when something is wrong, and take automatic actions when appropriate:

- *Amazon CloudWatch* monitors your AWS resources and the applications you run on AWS in real time. You can collect and track metrics, create customized dashboards, and set alarms that notify you or take actions when a specified metric reaches a threshold that you specify. For example, you can have CloudWatch track CPU usage or other metrics of your Amazon EC2 instances and automatically launch new instances when needed. For more information, see the [Amazon CloudWatch User Guide](#).
- *Amazon CloudWatch Events* delivers a near real-time stream of system events that describe changes in AWS resources. CloudWatch Events allows automated event-driven computing, as you can write rules that watch for certain events and trigger automated actions in other AWS services when these events happen. For more information, see the [Amazon CloudWatch Events User Guide](#).
- *Amazon CloudWatch Logs* allows you to monitor, store, and access your log files from Amazon EC2 instances, CloudTrail, and other sources. CloudWatch Logs can monitor information in the log files and notify you when certain thresholds are met. You can also archive your log data in highly durable storage. For more information, see the [Amazon CloudWatch Logs User Guide](#).
- *AWS CloudTrail* captures API calls and related events made by or on behalf of your AWS account and delivers the log files to an Amazon S3 bucket that you specify. You can identify which users and accounts called AWS, the source IP address from which the calls were made, and when the calls occurred. For more information, see the [AWS CloudTrail User Guide](#).

Monitoring Application Migration Service with Amazon CloudWatch

You can monitor AWS Transform MGN using CloudWatch, which collects raw data and processes it into readable, near real-time metrics. These statistics are kept for 15 months, so that you can access historical information and gain a better perspective on how your web application or service is performing. You can also set alarms that watch for certain thresholds, and send notifications or

take actions when those thresholds are met. For more information, see the [Amazon CloudWatch User Guide](#).

AWS Transform MGN supports six CloudWatch metrics in the AWS/MGN namespace.

AWS Transform MGN includes the following metrics across all Source servers. The following metrics are dimensionless.

Metric name	Description
ActiveSourceServerCount	Number of Source servers that are not archived.
TotalSourceServerCount	Number of source servers, including those that are archived.

AWS Transform MGN includes the following metrics by individual source server. The following metrics have a single dimension: **SourceServerID**.

Metric name	Description
LagDuration	The amount of time that has passed since the last consistent snapshot.
Backlog	The amount of data yet to be synced.
DurationSinceLastTest	The amount of time that has passed since the last Test instance launch.
ElapsedReplicationDuration	The cumulative amount of time this server has been replicating for (from which billing information is derived).

AWS Transform MGN EventBridge sample events

AWS Transform MGN sends events to Amazon EventBridge whenever a Source server launch has completed, a Source server reaches the READY_FOR_TEST lifecycle state for the first time, and when the data replication state becomes Stalled or when the data replication state is no

longer Stalled . You can use EventBridge and these events to write rules that take actions, such as notifying you, when a relevant event occurs. For more information, see [What is Amazon EventBridge?](#)

AWS Transform MGN sends events on a best-effort basis to EventBridge. Event delivery is not guaranteed.

Event samples

The following are sample MGN events in EventBridge:

Topics

- [MGN source server launch result](#)
- [MGN source server lifecycle state change](#)
- [MGN source server data replication stalled change](#)

MGN source server launch result

Emitted when a test or cutover instance launch was completed (successfully or with failure).

Possible states (referring to the **state** field within the **details** field):

1. TEST_LAUNCH_SUCCEEDED
2. TEST_LAUNCH_FAILED
3. CUTOVER_LAUNCH_SUCCEEDED
4. CUTOVER_LAUNCH_FAILED

Sample event:

```
{
  "version": "0",
  "id": "9da9af57-9253-4406-87cb-7cc400e43465",
  "detail-type": "MGN Source Server Launch Result",
  "source": "aws.mgn",
  "account": "111122223333",
  "time": "2016-08-22T20:12:19Z",
```

```

"region": "us-west-2",
"resources": [
  "arn:aws:mgn:us-west-2:111122223333:source-server/s-12345678901234567"
],
"detail": {
  "state": "*TEST_LAUNCH_SUCCEEDED*",      "job-id": "*mgnjob-04ca7d0d3fb6afa3e*"
}
}

```

MGN source server lifecycle state change

Emitted when a source server reaches the `READY_FOR_TEST` lifecycle state for the first time.

Sample event:

```

{
  "version": "0",
  "id": "9da9af57-9253-4406-87cb-7cc400e43465",
  "detail-type": "MGN Source Server Lifecycle State Change",
  "source": "aws.mgn",
  "account": "111122223333",
  "time": "2016-08-22T20:12:19Z",
  "region": "us-west-2",
  "resources": [
    "arn:aws:mgn:us-west-2:111122223333:source-server/s-12345678901234567"
  ],
  "detail": {
    "state": "*READY_FOR_TEST*"
  }
}

```

MGN source server data replication stalled change

Emitted when the data replication state becomes stalled, and when data replication state is no longer stalled (not stalled).

Possible states (referring to the **state** field within the **details** field):

1. STALLED

2. NOT_STALLED

Sample event:

```
{
  "version": "0",
  "id": "9da9af57-9253-4406-87cb-7cc400e43465",
  "detail-type": "MGN Source Server Data Replication Stalled Change",
  "source": "aws.mgn",
  "account": "111122223333",
  "time": "2016-08-22T20:12:19Z",
  "region": "us-west-2",
  "resources": [
    "arn:aws:mgn:us-west-2:111122223333:source-server/s-12345678901234567"
  ],
  "detail": {
    "state": "*STALLED*"
  }
}
```

Registering event rules

You create CloudWatch Events event rules that capture events coming from your AWS Transform MGN resources.

Note

When you use the AWS Management Console to create an event rule, the console automatically adds the IAM permissions necessary to grant Amazon CloudWatch Events permissions to call your desired target type. If you are creating an event rule using the AWS CLI, you must grant permissions explicitly. For more information, see [Events and Event Patterns](#) in the Amazon CloudWatch User Guide.

To create your CloudWatch Events rules:

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.

2. On the navigation pane, choose **Events, Create rule**.
3. For **Event source**, select **Event Pattern** as the event source, and then select **Build custom event pattern**.
4. Paste one of the following event patterns into the text area, depending on the event rule you wish to create:
 - a. To catch all MGN events:

```
{
  "source": ["aws.mgn"]
}
```

- b. To catch all Lifecycle state changes:

```
{
  "detail-type": ["MGN Source Server Lifecycle State Change"],
  "source": ["aws.mgn"]
}
```

- c. To catch all events relating to a given source server:

```
{
  "source": ["aws.mgn"],
  "resources": [
    "arn:aws:mgn:us-west-2:111122223333:source-server/s-12345678901234567"
  ]
}
```

5. For **Targets**, choose **Add target**. For **Target type**, choose your desired target.
6. Choose **Configure details**.
7. For **Rule definition**, type a name and description for your rule and choose **Create rule**.

Logging AWS Transform MGN with AWS CloudTrail

AWS Transform MGN is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service in AWS Transform MGN. CloudTrail captures all API calls for AWS Transform MGN as events. The calls captured include calls from the AWS Transform MGN console and code calls to the AWS Transform MGN API operations. If you create a trail, you can allow a continuous delivery of CloudTrail events to an Amazon S3 bucket, including events for

AWS Transform MGN. If you don't configure a trail, you can still view the most recent events in the CloudTrail console in **Event history**. Using the information collected by CloudTrail, you can determine the request that was made to AWS Transform MGN, the IP address from which the request was made, who made the request, when it was made, and additional details.

To learn more about CloudTrail, see the [AWS CloudTrail User Guide](#).

AWS Transform MGN information in CloudTrail

CloudTrail is activated on your AWS account when you create the account. When activity occurs in AWS Transform MGN, that activity is recorded in a CloudTrail event along with other AWS service events in **Event history**. You can view, search, and download recent events in your AWS account. For more information, see [Viewing events with CloudTrail Event history](#).

For an ongoing record of events in your AWS account, including events for AWS Transform MGN, create a trail. A *trail* allows CloudTrail to deliver log files to an Amazon S3 bucket. By default, when you create a trail in the console, the trail applies to all AWS Regions. The trail logs events from all Regions in the AWS partition and delivers the log files to the Amazon S3 bucket that you specify. Additionally, you can configure other AWS services to further analyze and act upon the event data collected in CloudTrail logs. For more information, see the following:

- [Overview for creating a trail](#)
- [CloudTrail supported services and integrations](#)
- [Configuring Amazon SNS notifications for CloudTrail](#)
- [Receiving CloudTrail log files from multiple regions](#) and [Receiving CloudTrail log files from multiple accounts](#)

All AWS Transform MGN actions are logged by CloudTrail and are documented in the AWS Transform MGN API. For example, calls to the `DescribeSourceServers` action generate entries in the CloudTrail log files.

Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root or AWS Identity and Access Management (IAM) user credentials.
- Whether the request was made with temporary security credentials for a role or federated user.

- Whether the request was made by another AWS service.

For more information, see the [CloudTrail userIdentity element](#).

Understanding AWS Transform MGN log file entries

A trail is a configuration that allows for the delivery of events as log files to an Amazon S3 bucket that you specify. CloudTrail log files contain one or more log entries. An event represents a single request from any source and includes information about the requested action, the date and time of the action, request parameters, and so on. CloudTrail log files aren't an ordered stack trace of the public API calls, so they don't appear in any specific order.

The following example shows a CloudTrail log entry that demonstrates the `DescribeSourceServers` action.

```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "AAAAAAAAAAAAAAAAAAAA",
    "arn": "arn:aws:sts::1234567890:assumed-role/Admin/user",
    "accountId": "1234567890",
    "accessKeyId": "BBBBBBBBBBBBBBBBBBBB",
    "sessionContext": {
      "sessionIssuer": {
        "type": "Role",
        "principalId": "AAAAAAAAAAAAAAAAAAAA",
        "arn": "arn:aws:iam::1234567890:role/Admin",
        "accountId": "1234567890",
        "userName": "Admin"
      },
      "webIdFederationData": {},
      "attributes": {
        "creationDate": "2021-10-20T14:19:17Z",
        "mfaAuthenticated": "false"
      }
    }
  },
  "eventTime": "2021-10-20T14:19:59Z",
  "eventSource": "mgn.amazonaws.com",
  "eventName": "DescribeSourceServers",
```

```
"awsRegion": "eu-west-1",
"sourceIPAddress": "54.240.197.234",
"userAgent": "Mozilla/5.0 (X11; Linux x86_64) AppleWebKit/537.36 (KHTML, like
Gecko) Chrome/94.0.4606.81 Safari/537.36",
"requestParameters": {
  "maxResults": 1000,
  "filters": {}
},
"responseElements": null,
"requestID": "d7618669-db08-4b53-bf6e-8a2cd57a677d",
"eventID": "436c17a7-3a54-4f4e-815d-4d980339744e",
"readOnly": true,
"eventType": "AwsApiCall",
"managementEvent": true,
"recipientAccountId": "1234567890",
"eventCategory": "Management"
}
```

Security in AWS Transform MGN

Topics

- [Overview](#)
- [Identity and access management for AWS Transform MGN](#)
- [Managing access using policies](#)
- [Using service-linked roles for AWS Transform MGN](#)
- [Policy structure](#)
- [Resilience in AWS Transform MGN](#)
- [Infrastructure security in AWS Transform MGN](#)
- [Compliance validation for AWS Transform MGN](#)
- [Cross-service confused deputy prevention](#)

Overview

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from a data center and network architecture that is built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between AWS and you. The [shared responsibility model](#) describes this as security of the cloud and security in the cloud:

- **Security of the cloud** – AWS is responsible for protecting the infrastructure that runs AWS services in the AWS Cloud. AWS also provides you with services that you can use securely. Third-party auditors regularly test and verify the effectiveness of our security as part of the [AWS Compliance Programs](#). To learn about the compliance programs that apply to AWS Transform MGN, see [AWS Services in Scope by Compliance Program](#).
- **Security in the cloud** – Your responsibility is determined by the AWS service that you use. You are also responsible for other factors including the sensitivity of your data, your company's requirements, and applicable laws and regulations

This documentation helps you understand how to apply the shared responsibility model when using AWS Transform MGN. It shows you how to configure AWS Transform MGN to meet your

security and compliance objectives. You also learn how to use other AWS services that help you to monitor and secure your AWS Transform MGN resources.

The customer is responsible for making sure that no misconfigurations are present during and after the migration process, including:

1. Access to replication servers should be allowed only from source servers CIDR range by applying proper security groups rules on replication servers.
2. After the migration, the customer should make sure that only allowed ports are exposed to the public internet.
3. Hardening of OS packages and other software deployed in the servers is completely under the customer's responsibility and we recommend the following:
 - a. Packages should be up to date and free of known vulnerabilities.
 - b. Only necessary OS/application services should be up and running.
4. Enabling the Anti-DDOS protection (AWS Shield) in the customer's AWS Account to eliminate the risk of denial of service attacks on the replication servers as well as the migrated servers.

Identity and access management for AWS Transform MGN

AWS Identity and Access Management (IAM) is an AWS service that helps an administrator securely control access to AWS resources. IAM administrators control who can be authenticated (signed in) and authorized (have permissions) to use AWS resources. IAM enables you to create users and groups under your AWS account. You control the permissions that users have to perform tasks using AWS resources. You can use IAM for no additional charge.

By default, users created via the IAM service don't have permissions for AWS Transform MGN resources and operations. To allow these users to manage AWS Transform MGN resources, you must create an IAM policy that explicitly grants them permissions, and attach the policy to the users or groups that require those permissions.

When you attach a policy to a user or group of users, it allows or denies the users permission to perform the specified tasks on the specified resources. For more information, see [Policies and Permissions](#) in the *IAM User Guide*.

Federated identity

As a best practice, require human users to use federation with an identity provider to access AWS services using temporary credentials.

A *federated identity* is a user from your enterprise directory, web identity provider, or Directory Service that accesses AWS services using credentials from an identity source. Federated identities assume roles that provide temporary credentials.

For centralized access management, we recommend AWS IAM Identity Center. For more information, see [What is IAM Identity Center?](#) in the *AWS IAM Identity Center User Guide*.

Authenticating with identities in AWS Transform MGN

Authentication is how you sign in to AWS using your identity credentials. You must be authenticated as the AWS account root user, an IAM user, or by assuming an IAM role.

You can sign in as a federated identity using credentials from an identity source like AWS IAM Identity Center (IAM Identity Center), single sign-on authentication, or Google/Facebook credentials. For more information about signing in, see [How to sign in to your AWS account](#) in the *AWS Sign-In User Guide*.

For programmatic access, AWS provides an SDK and CLI to cryptographically sign requests. For more information, see [AWS Signature Version 4 for API requests](#) in the *IAM User Guide*.

AWS account root user

When you create an AWS account, you begin with one sign-in identity called the AWS account *root user* that has complete access to all AWS services and resources. We strongly recommend that you don't use the root user for everyday tasks. For tasks that require root user credentials, see [Tasks that require root user credentials](#) in the *IAM User Guide*.

IAM users and groups

An [IAM user](#) is an identity with specific permissions for a single person or application. We recommend using temporary credentials instead of IAM users with long-term credentials. For more information, see [Require human users to use federation with an identity provider to access AWS using temporary credentials](#) in the *IAM User Guide*.

An [IAM group](#) specifies a collection of IAM users and makes permissions easier to manage for large sets of users. For more information, see [Use cases for IAM users](#) in the *IAM User Guide*.

IAM roles

An [IAM role](#) is an identity with specific permissions that provides temporary credentials. You can assume a role by [switching from a user to an IAM role \(console\)](#) or by calling an AWS CLI or AWS API operation. For more information, see [Methods to assume a role](#) in the *IAM User Guide*.

IAM roles are useful for federated user access, temporary IAM user permissions, cross-account access, cross-service access, and applications running on Amazon EC2. For more information, see [Cross account resource access in IAM](#) in the *IAM User Guide*.

Grant permission to tag resources during creation

Some resource-creating Amazon MGN API actions enable you to specify tags when you create the resource. You can use resource tags to implement attribute-based control (ABAC).

To enable users to tag resources on creation, they must have permissions to use the action that creates the resource, such as `mgn:RegisterAgentForMgn`. If tags are specified in the resource-creating action, Amazon performs additional authorization on the `mgn:TagResource` action to verify if users have permissions to create tags. Therefore, users must also have explicit permissions to use the `mgn:TagResource` action.

In the IAM policy definition for the `mgn:TagResource` action, use the `Condition` element with the `mgn:CreateAction` condition key to give tagging permissions to the action that creates the resource. The following example demonstrates a policy that allows an agent installer to create a source server and apply any tags to the source server on creation. The installer is not permitted to tag any existing resources (it cannot call the `mgn:TagResource` action directly).

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "mgn:SendAgentMetricsForMgn",
        "mgn:SendAgentLogsForMgn",
        "mgn:SendClientMetricsForMgn",
        "mgn:SendClientLogsForMgn"
      ]
    }
  ]
}
```



```

    ],
    "Resource": "*"
  },
  {
    "Effect": "Allow",
    "Action": [
      "mgn:RegisterAgentForMgn",
      "mgn:UpdateAgentSourcePropertiesForMgn",
      "mgn:UpdateAgentReplicationInfoForMgn",
      "mgn:UpdateAgentConversionInfoForMgn",
      "mgn:GetAgentInstallationAssetsForMgn",
      "mgn:GetAgentCommandForMgn",
      "mgn:GetAgentConfirmedResumeInfoForMgn",
      "mgn:GetAgentRuntimeConfigurationForMgn",
      "mgn:UpdateAgentBacklogForMgn",
      "mgn:GetAgentReplicationInfoForMgn"
    ],
    "Resource": "*"
  },
  {
    "Effect": "Allow",
    "Action": "mgn:TagResource",
    "Resource": "arn:aws:mgn:*:*:source-server/*",
    "Condition": {
      "StringEquals": {
        "mgn:CreateAction": "RegisterAgentForMgn"
      }
    }
  }
]
}

```

The `mgn:TagResource` action is only evaluated if tags are applied during the resource-creating action. Therefore, an installer that has permissions to create a resource (assuming there are no tagging conditions) does not require permissions to use the `mgn:TagResource` action if no tags are specified in the request. However, if the installer attempts to create a resource with tags, the request fails if the installer does not have permissions to use the `mgn:TagResource` action.

AWS managed policies for AWS Transform MGN

An AWS managed policy is a standalone policy that is created and administered by AWS. AWS managed policies are designed to provide permissions for many common use cases so that you can start assigning permissions to users, groups, and roles.

Keep in mind that AWS managed policies might not grant least-privilege permissions for your specific use cases because they're available for all AWS customers to use. We recommend that you reduce permissions further by defining [customer managed policies](#) that are specific to your use cases.

You cannot change the permissions defined in AWS managed policies. If AWS updates the permissions defined in an AWS managed policy, the update affects all principal identities (users, groups, and roles) that the policy is attached to. AWS is most likely to update an AWS managed policy when a new AWS service is launched or new API operations become available for existing services.

For more information, see [AWS managed policies](#) in the *IAM User Guide*.

MGN updates for AWS managed policies

View details about updates to AWS managed policies for AWS Transform MGN since March 1, 2021.

Change	Description	Date
AWSApplicationMigrationNetworkMigrationMultiAccount – Updated policy	Added <code>iam:CreateServiceLinkedRole</code> permission to create the <code>AWSServiceRoleForVPCTransitGateway</code> service-linked role, required for creation of Transit Gateway VPC attachments.	September 16, 2026
AWSApplicationMigrationNetworkMigrationMultiAccount – Updated policy	Added <code>cloudformation:TagResource</code> and <code>cloudformation:UntagResource</code> permissions for CloudFormation stacks	July 30, 2026

Change	Description	Date
	during create and update operations.	
AWSApplicationMigrationReadOnlyAccess – Updated policy	Added <code>ssm:GetParameter</code> permission to allow exporting post launch actions in the export file.	July 02, 2026
AWSApplicationMigrationNetworkMigrationMultiAccount – Updated policy	Added <code>ec2:DescribeAvailabilityZones</code> permission to resolve Availability Zone mappings in target member accounts.	June 28, 2026
AWSApplicationMigrationNetworkMigrationMultiAccount – Updated policy	Added <code>lambda:ListTags</code> permission for Lambda functions with prefix <code>[network-migration*]</code> .	June 8, 2026

Change	Description	Date
AWSApplicationMigrationFullAccess – Updated policy AWSApplicationMigrationReplicationServerPolicy – Updated policy AWSApplicationMigrationFSxProxyPolicy – New policy AWSApplicationMigrationFSxProxyVPCPolicy – New policy	Updated the AWSApplicationMigrationFullAccess policy to support FSx for ONTAP storage virtual machine operations. Updated the AWSApplicationMigrationReplicationServerPolicy policy to support snapshot group operations. Created new AWSApplicationMigrationFSxProxyPolicy policy to provide permissions to manage ONTAP file systems. Created new AWSApplicationMigrationFSxProxyVPCPolicy policy to provide permissions to manage PrivateLink between AWS Transform MGN and customer's FSx file system.	May 31, 2026
AWSApplicationMigrationNetworkMigrationCustomResource – New policy	MGN added a new policy that provides permissions for Network Migration custom resource. This identity-based policy allows modification of Transit Gateway resources that were specifically created by MGN.	December 2, 2025

Change	Description	Date
AWSApplicationMigrationNetworkMigrationMultiAccount – New policy	Created new managed policy to provide permissions to automate VMware to AWS network infrastructure migration through CloudFormation.	November 10, 2025
Updated the AWSApplicationMigrationSSMAccess and AWSApplicationMigrationFullAccess policies to support changes in SSM.	Added permission to tag network instances during RunInstances.	July 3, 2025
AWSApplicationMigrationServiceRolePolicy – Updated policy	Added permission to tag network instances during RunInstances.	March 13, 2025
AWSApplicationMigrationEC2Access – Updated policy	Added permission to tag network instances during RunInstances.	February 11, 2025
AWSApplicationMigrationServiceRolePolicy – Updated policy AWSApplicationMigrationEC2Access – Updated policy	Created new revisions of AWSApplicationMigrationServiceRolePolicy and AWSApplicationMigrationEC2Access managed policies to support a change in authentication with EBS APIs.	January 08, 2025
AWSApplicationMigrationFullAccess – Updated policy	Updated the AWSApplicationMigrationFullAccess policy to support SecureString parameter type in SSM Parameters Store for post-migration framework actions.	March 10, 2024

Change	Description	Date
AWSApplicationMigrationServiceEc2InstancePolicy – Updated policy	Created a new revision of the managed policy to support MGN in GovCloud and added SID to statements in the managed policy	December 28, 2023
AWSApplicationMigrationServiceEc2InstancePolicy – New policy	This policy allows installing and using the AWS Replication Agent, which is used by AWS Transform MGN to migrate source servers that run on EC2 (cross-Region or cross-AZ). An IAM role with this policy should be attached (as an EC2 Instance Profile) to the EC2 Instances.	August 21, 2023
AWSApplicationMigrationServiceRolePolicy – Updated policy	Updated the AWSApplicationMigrationServiceRolePolicy with Organizations permissions to support the global view feature.	June 18, 2023
AWSApplicationMigrationFullAccess – Updated policy	Updated the AWSApplicationMigrationFullAccess policy to support specific automation SSM documents.	April 1, 2023

Change	Description	Date
AWSApplicationMigrationFullAccess – Updated policy AWSApplicationMigrationSSMAccess – Updated policy AWSApplicationMigrationReadOnlyAccess – Updated policy	Updated the AWSApplicationMigrationFullAccess policy to support both command and automation SSM documents for post-migration framework actions. Updated the AWSApplicationMigrationSSMAccess policy to support both command and automation SSM documents for the custom actions feature. Updated the AWSApplicationMigrationReadOnlyAccess policy to support the new import and export feature.	March 21, 2023
AWSApplicationMigrationEC2Access – Updated policy	Updated the AWSApplicationMigrationEC2Access policy to support: DescribeSnapshots, DescribeImages, DescribeVolumes.	January 29, 2023

Change	Description	Date
AWSApplicationMigrationEC2Access – Updated policy AWSApplicationMigrationReadOnlyAccess – Updated policy AWSApplicationMigrationSSMAccess – Created policy	Updated the AWSApplicationMigrationEC2Access policy to support: CreateLaunchTemplate, DeleteLaunchTemplate. Updated the AWSApplicationMigrationReadOnlyAccess policy to support: DescribeLaunchConfigurationTemplates, ListSourceServerActions, ListTemplateActions, ListApplications, ListWaves. Created new AWSApplicationMigrationSSMAccess policy to support new custom actions feature.	November 28, 2022
AWSApplicationMigrationAgentPolicy – Updated policy AWSApplicationMigrationAgentInstallationPolicy – Updated policy	Updated the AWSApplicationMigrationAgentPolicy policy and the AWSApplicationMigrationAgentInstallationPolicy policy to support sending additional metrics during the agent installation process.	September 20, 2022

Change	Description	Date
AWSApplicationMigrationAgentInstallationPolicy – New policy	MGN added a new policy. This policy allows installing the AWS Replication Agent, which is used with Application Migration Service to migrate source servers to AWS. Attach this policy to your users or roles whose credentials you provide during the installation step of the AWS Replication Agent. The installed AWS Replication Agent will communicate with AWS Transform MGN using the recommended strong authentication method.	June 15, 2022
AWSApplicationMigrationFullAccess – Updated policy	Updated the AWSApplicationMigrationFullAccess policy to support the Post Migration Framework.	May 16, 2022
AWSApplicationMigrationAgentPolicy_v2 – New policy	AWS Transform MGN added a new policy. This policy allows using the AWS Replication Agent, which is used with AWS Transform MGN to migrate source servers to AWS. We do not recommend that you attach this policy to your users or roles.	May 10, 2022

Change	Description	Date
AWSApplicationMigrationReadOnlyAccess – Updated policy	Updated the AWSApplicationMigrationReadOnlyAccess policy to include service quotas.	April 3, 2022
AWSApplicationMigrationEC2Access – Updated policy	Updated the AWSApplicationMigrationEC2Access policy to add additional permissions and restrict certain existing permissions. This policy is only intended to be used for the MGN console. The restriction prevents certain requests from being called directly by the calling identity, while enabling AWS Transform MGN to make the request to EC2 on behalf of the calling identity.	March 2, 2022
AWSApplicationMigrationServiceRolePolicy – Updated policy	AWS Transform MGN added a new policy to allow AWS Transform MGN to manage AWS resources on your behalf.	December 15, 2021
AWSApplicationMigrationVCenterClientPolicy – New policy	AWS Transform MGN added a new policy that allows the installation and usage of the AWS vCenter Appliance.	November 7, 2021
AWSApplicationMigrationAgentPolicy – New policy	AWS Transform MGN added a new policy to allow the installation of the AWS Replication Agent on source servers.	April 18, 2021

Change	Description	Date
AWSApplicationMigrationConversionServerPolicy – New policy	AWS Transform MGN added a new policy that allows AWS Transform MGN to communicate with the service.	April 18, 2021
AWSApplicationMigrationMGHAccess – New policy	AWS Transform MGN added a new policy to allow AWS Transform MGN access to your account's AWS Migration Hub	April 18, 2021
AWSApplicationMigrationReplicationServerPolicy – New policy	AWS Transform MGN added a new policy to allow the AWS Transform MGN replication servers to communicate with the service, create and manage resources on your behalf.	April 7, 2021
MGN started tracking changes	AWS Transform MGN started tracking changes for AWS managed policies.	April 7, 2021

Topics

- [AWS managed policy: AWSApplicationMigrationServiceRolePolicy](#)
- [AWS managed policy: AWSApplicationMigrationConversionServerPolicy](#)
- [AWS managed policy: AWSApplicationMigrationReplicationServerPolicy](#)
- [AWS managed policy: AWSApplicationMigrationFSxProxyPolicy](#)
- [AWS managed policy: AWSApplicationMigrationFSxProxyVPCPolicy](#)
- [AWS managed policy: AWSApplicationMigrationAgentPolicy](#)
- [AWS managed policy: AWSApplicationMigrationMGHAccess](#)
- [AWS managed policy: AWSApplicationMigrationFullAccess](#)

- [AWS managed policy: AWSApplicationMigrationEC2Access](#)
- [AWS managed policy: AWSApplicationMigrationSSMAccess](#)
- [AWS managed policy: AWSApplicationMigrationReadOnlyAccess](#)
- [AWS managed policy: AWSApplicationMigrationVCenterClientPolicy](#)
- [AWS managed policy: AWSApplicationMigrationAgentInstallationPolicy](#)
- [AWS managed policy: AWSApplicationMigrationAgentPolicy_v2](#)
- [AWS managed policy: AWSApplicationMigrationServiceEc2InstancePolicy](#)
- [AWS managed policy: AWSApplicationMigrationNetworkMigrationMultiAccount](#)
- [AWS managed policy: AWSApplicationMigrationNetworkMigrationCustomResource](#)

AWS managed policy: AWSApplicationMigrationServiceRolePolicy

This policy is attached to the MGN [AWSServiceRoleForApplicationMigrationService](#) service-linked role (SLR).

This policy allows AWS Transform MGN to manage AWS resources on your behalf.

Permissions details

To view the policy permission details see [AWSApplicationMigrationServiceRolePolicy](#) in the AWS Managed Policy Reference Guide.

AWS managed policy: AWSApplicationMigrationConversionServerPolicy

This policy is attached to the AWS Transform MGN conversion server's instance role.

This policy allows the AWS Transform MGN conversion servers, which are EC2 instances launched by AWS Transform MGN, to communicate with the MGN service. An IAM role with this policy is attached (as an EC2 Instance Profile) by MGN to the MGN Conversion Servers, which are automatically launched and terminated by MGN, when needed. We do not recommend that you attach this policy to your users or roles. MGN conversion servers are used by AWS Transform MGN when users choose to launch test or cutover instances using the MGN console, CLI, or API.

Permissions details

To view the policy permission details see [AWSApplicationMigrationConversionServerPolicy](#) in the AWS Managed Policy Reference Guide.

AWS managed policy: AWSApplicationMigrationReplicationServerPolicy

This policy is attached to the AWS Transform MGN replication server's instance role.

This policy allows the AWS Transform MGN Replication Servers, which are EC2 instances launched by AWS Transform MGN - to communicate with the AWS MGN service, and to create EBS snapshots in your AWS account. An IAM role with this policy is attached (as an EC2 Instance Profile) by AWS Transform MGN to the MGN replication servers which are automatically launched and terminated by MGN, as needed. MGN Replication Servers are used to help data replication from your external servers to AWS, as part of the migration process managed using MGN. We do not recommend that you attach this policy to your users or roles.

Permissions details

To view the policy permission details see [AWSApplicationMigrationReplicationServerPolicy](#) in the AWS Managed Policy Reference Guide.

AWS managed policy: AWSApplicationMigrationFSxProxyPolicy

This policy is attached to the AWS Transform MGN FSx proxy service role.

This policy provides permissions to manage ONTAP file systems as part of the AWS Transform MGN file system migration process. The policy allows the service to read secrets, describe and create FSx volumes, and manage tags on FSx resources. We do not recommend that you attach this policy to your users or roles.

Permissions details

To view the policy permission details see [AWSApplicationMigrationFSxProxyPolicy](#) in the AWS Managed Policy Reference Guide.

AWS managed policy: AWSApplicationMigrationFSxProxyVPCPolicy

This policy is attached to the AWS Transform MGN FSx proxy VPC service role.

This policy provides permissions to manage PrivateLink connectivity between AWS Transform MGN and customer's FSx file system. The policy allows the service to create and manage security groups,

VPC endpoint service configurations, Network Load Balancers, and target groups required for the PrivateLink connection. We do not recommend that you attach this policy to your users or roles.

Permissions details

To view the policy permission details see [AWSApplicationMigrationFSxProxyVPCPolicy](#) in the AWS Managed Policy Reference Guide.

AWS managed policy: AWSApplicationMigrationAgentPolicy

You can attach the AWSApplicationMigrationAgentPolicy policy to your IAM identities.

This policy allows installing and using the AWS Replication Agent, which is used with AWS Transform MGN to migrate external servers to AWS. Attach this policy to your users whose credentials you provide when installing the AWS replication agent.

Permissions details

To view the policy permission details see [AWSApplicationMigrationAgentPolicy](#) in the AWS Managed Policy Reference Guide.

AWS managed policy: AWSApplicationMigrationMGHAccess

This policy allows AWS Transform MGN to send metadata about the progress of servers being migrated using MGN to AWS Migration Hub (MGH). MGN automatically creates an IAM role with this policy attached and assumes this role. We do not recommend that you attach this policy to your users or roles. Migration-progress data is only sent after the AWS "home region" is set in AWS MGH. If the Home AWS Region is different than the AWS Region into which a server is being migrated, this data will be sent cross-region. To stop MGN from sending this metadata to AWS MGH, detach it from your users or roles.

Permissions details

To view the policy permission details see [AWSApplicationMigrationMGHAccess](#) in the AWS Managed Policy Reference Guide.

AWS managed policy: AWSApplicationMigrationFullAccess

You can attach the AWSApplicationMigrationFullAccess policy to your IAM identities.

This policy provides permissions to all public APIs of AWS Transform MGN, as well as permissions to read KMS key, License Manager, Resource Groups, Elastic Load Balancing, IAM, EC2, and Amazon FSx information. This policy should only be granted to an administrator or a power-user.

Important

You must attach the [AWSApplicationMigrationFullAccess](#) and the [AWSApplicationMigrationEC2Access](#) policies to your users and roles to enable them to launch test and cutover instances and to complete a full migration cycle with MGN.

Permissions details

To view the policy permission details see [AWSApplicationMigrationFullAccess](#) in the AWS Managed Policy Reference Guide.

AWS managed policy: AWSApplicationMigrationEC2Access

You can attach the AWSApplicationMigrationEC2Access policy to your IAM identities.

This policy allows Amazon EC2 operations required to use AWS Transform MGN to launch the migrated servers as EC2 instances. Attach this policy to your users or roles. This policy is only intended to be used for the MGN console.

Permissions details

To view the policy permission details see [AWSApplicationMigrationEC2Access](#) in the AWS Managed Policy Reference Guide.

AWS managed policy: AWSApplicationMigrationSSMAccess

You can attach the AWSApplicationMigrationSSMAccess policy to your IAM identities.

This policy allows Amazon SSM operations required to use AWS Transform MGN to run SSM documents post migration of source servers. Attach this policy to your users or roles. This policy is only intended to be used for the MGN console.

Permissions details

To view the policy permission details see [AWSApplicationMigrationSSMAccess](#) in the AWS Managed Policy Reference Guide.

AWS managed policy: AWSApplicationMigrationReadOnlyAccess

You can attach the `AWSApplicationMigrationReadOnlyAccess` policy to your IAM identities.

This policy provides permissions to all read-only public APIs of AWS Transform MGN, as well as some read-only APIs of other AWS services that are required in order to make full read-only use of the MGN console. It does not allow them to perform any actions, such as initialize the service, replicate servers, or launch servers in AWS. This policy can be granted to a user in a support role.

Attach this policy to your users or roles.

Permissions details

To view the policy permission details see [AWSApplicationMigrationReadOnlyAccess](#) in the AWS Managed Policy Reference Guide.

AWS managed policy: AWSApplicationMigrationVCenterClientPolicy

You can attach the `AWSApplicationMigrationVCenterClientPolicy` policy to your IAM identities.

This policy allows installing and using the MGN vCenter Client, which is used with AWS Transform MGN to migrate external servers to AWS. Attach this policy to your users or roles whose credentials you provide when installing the MGN vCenter Client.

Permissions details

To view the policy permission details see [AWSApplicationMigrationVCenterClientPolicy](#) in the AWS Managed Policy Reference Guide.

AWS managed policy: AWSApplicationMigrationAgentInstallationPolicy

This policy allows installing the AWS Replication Agent, which is used with AWS Transform MGN to migrate source servers to AWS. Attach this policy to your users or roles whose credentials you provide during the installation step of the AWS Replication Agent. The installed AWS Replication Agent will communicate with MGN using the recommended strong authentication method.

Permissions details

To view the policy permission details see [AWSApplicationMigrationAgentInstallationPolicy](#) in the AWS Managed Policy Reference Guide.

AWS managed policy: AWSApplicationMigrationAgentPolicy_v2

This policy allows using the AWS Replication Agent, which is used with AWS Transform MGN to migrate source servers to AWS. We do not recommend that you attach this policy to your users or roles.

Permissions details

To view the policy permission details see [AWSApplicationMigrationAgentPolicy_v2](#) in the AWS Managed Policy Reference Guide.

AWS managed policy: AWSApplicationMigrationServiceEc2InstancePolicy

This policy allows installing and using the AWS Replication Agent, which is used by AWS Transform MGN to migrate source servers that run on EC2 (cross-Region or cross-AZ). An IAM role with this policy should be attached (as an EC2 Instance Profile) to the EC2 Instances. This policy can only be used within the same AWS account.

Permissions details

To view the policy permission details see [AWSApplicationMigrationServiceEc2InstancePolicy](#) in the AWS Managed Policy Reference Guide.

AWS managed policy: AWSApplicationMigrationNetworkMigrationMultiAccount

You can attach the AWSApplicationMigrationNetworkMigrationMultiAccount policy to your IAM identities.

This identity-based policy enables AWS Transform MGN to create, modify, and manage network infrastructure components through CloudFormation. The policy grants permissions necessary for:

1. **Network Resource Management:** Creating and managing VPCs, subnets, route tables, and network ACLs; configuring Transit Gateways and their attachments; managing security groups and their rules; setting up NAT Gateways and Internet Gateways; handling network interfaces and elastic IPs
2. **CloudFormation Operations:** Creating and managing stacks with prefix [Nmd*]; describing stack resources and events; updating and deleting stacks
3. **Resource Sharing:** Managing RAM (Resource Access Manager) resource shares; sharing Transit Gateways across accounts within the same organization
4. **Custom Resources:** Creating and managing Lambda functions with prefix [network-migration*]; managing IAM roles with prefix [Nmd*modifyTransitGateway*]; creating and managing CloudWatch log groups

The policy enforces security through resource tagging requirements (CreatedBy: AWSApplicationMigrationService), conditional checks ensuring operations are called via CloudFormation, organization-level controls for cross-account resource sharing, and specific resource-level permissions for critical network components.

This policy grants both programmatic and console access required for AWS Transform MGN to orchestrate network infrastructure deployment and management through CloudFormation.

Permissions details

To view the policy permission details see [AWSApplicationMigrationNetworkMigrationMultiAccount](#) in the AWS Managed Policy Reference Guide.

AWS managed policy:

AWSApplicationMigrationNetworkMigrationCustomResource

Allows modification of Transit Gateway resources created by MGN. You can attach the AWSApplicationMigrationNetworkMigrationCustomResource policy to your IAM identities.

This identity-based policy allows modification of Transit Gateway resources that were specifically created by MGN. The policy grants permission to modify Transit Gateways and their route tables, but only if they are tagged with [CreatedBy: AWSApplicationMigrationService]. This restriction ensures that only resources created by the migration service can be modified, providing

targeted control over Transit Gateway infrastructure during migration processes. The policy grants the permissions necessary to complete these actions programmatically from the AWS API or AWS CLI.

The policy is particularly useful for:

- Managing Transit Gateway configurations during application migration
- Ensuring only migration service-created resources can be modified
- Maintaining control over network infrastructure changes during migration processes

Permissions details

To view the policy permission details see

[AWSApplicationMigrationNetworkMigrationCustomResource](#) in the AWS Managed Policy Reference Guide.

Managing access using policies

You control access in AWS by creating policies and attaching them to AWS identities or resources. A policy defines permissions when associated with an identity or resource. AWS evaluates these policies when a principal makes a request. Most policies are stored in AWS as JSON documents. For more information about JSON policy documents, see [Overview of JSON policies](#) in the *IAM User Guide*.

Using policies, administrators specify who has access to what by defining which **principal** can perform **actions** on what **resources**, and under what **conditions**.

By default, users and roles have no permissions. An IAM administrator creates IAM policies and adds them to roles, which users can then assume. IAM policies define permissions regardless of the method used to perform the operation.

Identity-based policies

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as a user, role, or group. These policies control what actions that identity can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see [Define custom IAM permissions with customer managed policies](#) in the IAM User Guide.

Identity-based policies can be further categorized as inline policies or managed policies. Inline policies are embedded directly into a single user, group, or role. Managed policies are standalone policies that you can attach to multiple users, groups, and roles in your AWS account. Managed policies include AWS managed policies and customer managed policies. To learn how to choose between a managed policy or an inline policy, see [Managed policies and inline policies](#) in the IAM User Guide.

Using identity-based policies

By default, users and roles don't have permission to create or modify AWS Transform MGN resources. They also can't perform tasks using the AWS Management Console, AWS CLI, or AWS API. An IAM administrator must create IAM policies that grant users and roles permission to perform specific API operations on the specified resources they need. The administrator must then attach those policies to the users or groups that require those permissions. To understand how to attach policies to a user or group, learn about [adding and removing IAM identity permissions](#). To learn how to create an IAM identity-based policy using example JSON policy documents, see [Creating policies on the JSON tab in the IAM User Guide](#).

Topics

- [Customer-managed policies in MGN](#)
- [Network Migration API permissions](#)

Customer-managed policies in MGN

You can create your own custom IAM policies to allow permissions for AWS Transform MGN actions and resources. You can attach these custom policies to the users, roles, or groups that require those permissions. You can also create your own custom IAM policies for integration between AWS Transform MGN and other AWS services. The next few topics provide examples of IAM policies that grant permission for various AWS Transform MGN actions. Use them to limit AWS Transform MGN access for your users and roles.

Network Migration API permissions

The Network Migration APIs allow you to automate the migration of network infrastructure from VMware to AWS. To use these APIs, attach both the [AWSApplicationMigrationNetworkMigrationMultiAccount](#) managed policy and the following custom policy to your IAM identity.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Tags",
      "Effect": "Allow",
      "Action": [
        "mgn:TagResource"
      ],
      "Resource": [
        "arn:aws:mgn:*:*:network-migration-definition/*"
      ],
      "Condition": {
        "StringEquals": {
          "aws:ResourceTag/CreatedBy": "AWSTransform",
          "mgn:CreateAction": [
            "CreateNetworkMigrationDefinition"
          ]
        }
      }
    },
    {
      "Sid": "CreateMethod",
      "Effect": "Allow",
      "Action": [
        "mgn:CreateNetworkMigrationDefinition"
      ],
      "Resource": [
        "*"
      ],
      "Condition": {
        "StringEquals": {
          "aws:RequestTag/CreatedBy": "AWSTransform"
        }
      }
    },
    {
      "Sid": "ResourceMethods",
      "Effect": "Allow",
      "Action": [
        "mgn:UpdateNetworkMigrationDefinition",

```

```

        "mgn:StartNetworkMigrationMapping",
        "mgn:StartNetworkMigrationCodeGeneration",
        "mgn:StartNetworkMigrationDeployment",
        "mgn:StartNetworkMigrationAnalysis"
    ],
    "Resource": [
        "*"
    ],
    "Condition": {
        "StringEquals": {
            "aws:ResourceTag/CreatedBy": "AWSTransform"
        }
    }
},
{
    "Sid": "ReadonlyMethods",
    "Effect": "Allow",
    "Action": [
        "mgn:GetNetworkMigrationDefinition"
    ],
    "Resource": [
        "arn:aws:mgn:*:*:network-migration-definition/*"
    ],
    "Condition": {
        "StringEquals": {
            "aws:ResourceTag/CreatedBy": "AWSTransform"
        }
    }
},
{
    "Sid": "DeleteExistingNetworkMigrationDefinition",
    "Effect": "Allow",
    "Action": [
        "mgn:DeleteNetworkMigrationDefinition"
    ],
    "Resource": [
        "arn:aws:mgn:*:*:network-migration-definition/*"
    ],
    "Condition": {
        "StringEquals": {
            "aws:ResourceTag/CreatedBy": "AWSTransform"
        }
    }
},

```

```

{
  "Sid": "ReadOnly",
  "Effect": "Allow",
  "Action": [
    "mgn:ListNetworkMigrationDefinitions",
    "mgn:ListNetworkMigrationExecutions",
    "mgn:ListNetworkMigrationMapperSegments",
    "mgn:ListNetworkMigrationMappings",
    "mgn:ListNetworkMigrationMapperSegmentConstructs",
    "mgn:ListNetworkMigrationCodeGenerationSegments",
    "mgn:ListNetworkMigrationCodeGenerations",
    "mgn:ListNetworkMigrationDeployedStacks",
    "mgn:ListNetworkMigrationDeployments",
    "mgn:ListNetworkMigrationAnalysisResults",
    "mgn:ListNetworkMigrationAnalyses",
    "mgn:GetNetworkMigrationMapperSegmentConstruct"
  ],
  "Resource": [
    "*"
  ]
},
{
  "Sid": "MGNNetworkMigrationUpdate",
  "Effect": "Allow",
  "Action": [
    "mgn:UpdateNetworkMigrationMapperSegment",
    "mgn:StartNetworkMigrationMappingUpdate",
    "mgn:ListNetworkMigrationMappingUpdates"
  ],
  "Resource": [
    "arn:aws:mgn:*:*:network-migration-definition/*"
  ],
  "Condition": {
    "StringEquals": {
      "aws:ResourceTag/CreatedBy": "AWSTransform"
    }
  }
},
{
  "Sid": "MGNImportFileEnrichment",
  "Effect": "Allow",
  "Action": [
    "mgn:StartImportFileEnrichment",
    "mgn:ListImportFileEnrichments"
  ]
}

```

```

    ],
    "Resource": [
        "*"
    ]
},
{
    "Sid": "S3Bucket",
    "Effect": "Allow",
    "Action": [
        "s3:ListBucket",
        "s3:GetBucketTagging",
        "s3:GetBucketPublicAccessBlock",
        "s3:GetBucketLocation",
        "s3:CreateBucket",
        "s3:PutBucketTagging",
        "s3:PutEncryptionConfiguration"
    ],
    "Resource": "arn:aws:s3:::*",
    "Condition": {
        "ForAnyValue:StringEquals": {
            "aws:CalledVia": [
                "mgn.amazonaws.com"
            ]
        }
    }
},
{
    "Sid": "S3BucketObject",
    "Effect": "Allow",
    "Action": [
        "s3:GetObject",
        "s3:GetObjectVersion",
        "s3:ListMultipartUploadParts",
        "s3:ListBucketMultipartUploads",
        "s3:GetObjectAttributes",
        "s3:PutObject",
        "s3:AbortMultipartUpload",
        "s3:DeleteObject"
    ],
    "Resource": "arn:aws:s3::*/*",
    "Condition": {
        "ForAnyValue:StringEquals": {
            "aws:CalledVia": [
                "mgn.amazonaws.com"
            ]
        }
    }
}

```



```

        ]
      }
    }
  },
  {
    "Sid": "MGNNetworkAnalysis",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInsightsPath",
      "ec2:StartNetworkInsightsAnalysis",
      "ec2>DeleteNetworkInsightsPath",
      "ec2>DeleteNetworkInsightsAnalysis",
      "ec2:CreateTags"
    ],
    "Resource": [
      "arn:aws:ec2:*:*:network-insights-path/*",
      "arn:aws:ec2:*:*:network-insights-analysis/*",
      "arn:aws:ec2:*:*:network-interface/*"
    ],
    "Condition": {
      "ForAnyValue:StringEquals": {
        "aws:CalledVia": [
          "mgn.amazonaws.com"
        ]
      }
    }
  },
  {
    "Sid": "EC2DescribeNoCondition",
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeVpcAttribute"
    ],
    "Resource": "*"
  },
  {
    "Sid": "MGNServiceQuota",
    "Effect": "Allow",
    "Action": "servicequotas:GetServiceQuota",
    "Resource": "arn:aws:servicequotas:*:*:vpc/L-2AFB9258",
    "Condition": {
      "ForAnyValue:StringEquals": {
        "aws:CalledVia": "mgn.amazonaws.com"
      }
    }
  }
}

```

```

    }
  },
  {
    "Sid": "EC2GetSubnetCidrReservations",
    "Effect": "Allow",
    "Action": "ec2:GetSubnetCidrReservations",
    "Resource": "*"
  },
  {
    "Sid": "TiroForNetworkInsights",
    "Effect": "Allow",
    "Action": [
      "tiros:CreateQuery",
      "tiros:GetQueryAnswer",
      "tiros:GetQueryExplanation"
    ],
    "Resource": "*"
  }
]
}

```

Resource-based policies

Resource-based policies are JSON policy documents that you attach to a resource. Examples include IAM *role trust policies* and Amazon S3 *bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. You must [specify a principal](#) in a resource-based policy.

Resource-based policies are inline policies that are located in that service. You can't use AWS managed policies from IAM in a resource-based policy.

Access control lists (ACLs)

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

Amazon S3, AWS WAF, and Amazon VPC are examples of services that support ACLs. To learn more about ACLs, see [Access control list \(ACL\) overview](#) in the *Amazon Simple Storage Service Developer Guide*.

Other policy types

AWS supports additional policy types that can set the maximum permissions granted by more common policy types:

- **Permissions boundaries** – Set the maximum permissions that an identity-based policy can grant to an IAM entity. For more information, see [Permissions boundaries for IAM entities](#) in the *IAM User Guide*.
- **Service control policies (SCPs)** – Specify the maximum permissions for an organization or organizational unit in AWS Organizations. For more information, see [Service control policies](#) in the *AWS Organizations User Guide*.
- **Resource control policies (RCPs)** – Set the maximum available permissions for resources in your accounts. For more information, see [Resource control policies \(RCPs\)](#) in the *AWS Organizations User Guide*.
- **Session policies** – Advanced policies passed as a parameter when creating a temporary session for a role or federated user. For more information, see [Session policies](#) in the *IAM User Guide*.

Multiple policy types

When multiple types of policies apply to a request, the resulting permissions are more complicated to understand. To learn how AWS determines whether to allow a request when multiple policy types are involved, see [Policy evaluation logic](#) in the *IAM User Guide*.

Using service-linked roles for AWS Transform MGN

AWS Transform MGN uses AWS Identity and Access Management (IAM) [service-linked roles](#). A service-linked role is a unique type of IAM role that is linked directly to AWS Transform MGN. Service-linked roles are predefined by AWS Transform MGN and include all the permissions that the service requires to call other AWS services on your behalf.

A service-linked role makes setting up AWS Transform MGN easier because you don't have to manually add the necessary permissions. AWS Transform MGN defines the permissions of its service-linked roles, and unless defined otherwise, only AWS Transform MGN can assume its roles. The defined permissions include the trust policy and the permissions policy, and that permissions policy cannot be attached to any other IAM entity.

You can delete a service-linked role only after first deleting their related resources. This protects your AWS Transform MGN resources because you can't inadvertently remove permission to access the resources.

For information about other services that support service-linked roles, see [AWS Services That Work with IAM](#) and look for the services that have **Yes** in the **Service-Linked Role** column. Choose a **Yes** with a link to view the service-linked role documentation for that service.

AWSServiceRoleForApplicationMigrationService service-linked role

AWS Transform MGN uses the service-linked role named **AWSServiceRoleForApplicationMigrationService**. This service-linked role has scoped permissions that AWS Transform MGN needs to run in your account. Its permissions are defined by the `AWSApplicationMigrationServiceRolePolicy` AWS managed policy.

The `AWSServiceRoleForApplicationMigrationService` service-linked role trusts the `mgn.amazonaws.com` service principal to assume the role. The role permissions are defined in the [AWSApplicationMigrationServiceRolePolicy](#) AWS managed policy.

To view the policy permission details see [AWSApplicationMigrationServiceRolePolicy](#) in the AWS Managed Policy Reference Guide.

You must configure permissions to allow an IAM entity (such as a user, group, or role) to create, edit, or delete a service-linked role. For more information, see [Service-Linked Role Permissions](#) in the *IAM User Guide*.

Creating a service-linked role for AWS Transform MGN

You don't need to manually create a service-linked role. When you configure the Replication Configuration Template for AWS Transform MGN, a service-linked role is automatically created. MGN automatically creates the IAM service-linked role, which you can see in the IAM console. You don't need to manually create or configure this role.

If you delete this service-linked role, and then need to create it again, you can use the same process to recreate the role in your account. When you create the first new replication configuration template in MGN, it creates the service-linked role for you again.

In the AWS CLI or the AWS API, create a service-linked role with the AWS Transform MGN name. For more information, see [Creating a Service-Linked Role](#) in the *IAM User Guide*. If you delete this service-linked role, you can use this same process to create the role again.

Editing a service-linked role for AWS Transform MGN

AWS Transform MGN does not allow you to edit the `AWSServiceRoleForApplicationMigrationService` service-linked role. After you create a service-linked role, you cannot change the name of the role because various entities might reference the role. However, you can edit the description of the role using IAM. For more information, see [Editing a Service-Linked Role](#) in the *IAM User Guide*.

Deleting a service-linked role for AWS Transform MGN

If you no longer need to use a feature or service that requires a service-linked role, we recommend that you delete that role. That way you don't have an unused entity that is not actively monitored or maintained. However, you must clean up the resources for your service-linked role before you can manually delete it.

Note

If AWS Transform MGN is using the role when you try to delete the resources, the deletion might fail. If that happens, wait for a few minutes and try the operation again.

Warning

Resources cannot be cleaned up without stopping the services provided by AWS Transform MGN. Cleaning up these resources will cause AWS Transform MGN to stop working. Before you run the following commands, confirm that migrations and replication are no longer needed. For more information, see [Cleaning up a Service-Linked Role](#) in the *IAM User Guide*.

To clean up AWS Transform MGN resources used by `AWSServiceRoleForApplicationMigrationService`

1. Identify and delete any waves and applications in all AWS Regions
 - a. identify any waves:

```
aws mgn list-waves
```

b. Delete any waves:

```
aws mgn delete-wave --wave-id {WaveID}
```

c. Identify any application:

```
aws mgn list-applications
```

d. Delete any application:

```
aws mgn delete-application --application-id {ApplicationID}
```

2. Identify and delete any source servers in all AWS Regions**a. Identify any active source servers:**

```
aws mgn describe-source-servers --filters isArchived=False --query  
"items[*].sourceServerID"
```

b. Disconnect any active source servers:

```
aws mgn disconnect-from-service --source-server-id {SourceServerID}
```

c. Archive any disconnected source servers:

```
aws mgn mark-as-archived --source-server-id {SourceServerID}
```

d. Delete any archived source server:

```
aws mgn delete-source-server --source-server-id {SourceServerID}
```

3. Identify and delete any MGN jobs in all AWS Regions**a. Identify any MGN jobs**

```
aws mgn describe-jobs
```

b. Delete any MGN jobs:

```
aws mgn delete-job --job-id {MGNJobId}
```

4. Identify and delete any MGN replication templates

a. Identify any MGN replication template:

```
aws mgn describe-replication-configuration-templates
```

b. Remove any MGN replication templates:

```
aws mgn delete-replication-configuration-template --replication-configuration-  
template-id {rct-TemplateID}
```

To manually delete the service-linked role using IAM

Use the IAM console, the AWS CLI, or the AWS API to delete the `AWSServiceRoleForApplicationMigrationService` service-linked role. For more information, see [Deleting a Service-Linked Role](#) in the *IAM User Guide*.

Supported Regions for MGN service-linked roles

AWS Transform MGN supports using service-linked roles in all of the [AWS Regions where the service is available](#).

Policy structure

An IAM policy is a JSON document that consists of one or more statements. Each statement is structured as follows.

```
{  
  "Statement": [  
    {  
      "Effect": "effect",  
      "Action": "action",  
      "Resource": "arn",  
      "Condition": {  
        "condition": {  
          "key": "value"  
        }  
      }  
    }  
  ]  
}
```

There are various elements that make up a statement:

- **Effect:** The effect can be Allow or Deny. By default, IAM users don't have permission to use resources and API actions, so all requests are denied. An explicit allow overrides the default. An explicit deny overrides any allows.
- **Action:** The action is the specific AWS Transform MGN API action for which you are granting or denying permission.
- **Resource:** The resource that's affected by the action. For AWS Transform MGN, you must specify "*" as the resource.
- **Condition:** Conditions are optional. They can be used to control when your policy is in effect.

Resilience in AWS Transform MGN

The AWS global infrastructure is built around AWS Regions and Availability Zones. Regions provide multiple physically separated and isolated Availability Zones, which are connected through low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that automatically fail over between zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

For more information about AWS Regions and Availability Zones, see [AWS Global Infrastructure](#).

Infrastructure security in AWS Transform MGN

As a managed service, AWS Transform MGN is protected by the AWS global network security procedures that are described in the [Amazon Web Services: Overview of Security Processes](#) whitepaper.

You use AWS published API calls to access AWS Transform MGN through the network. Clients must support Transport Layer Security (TLS) 1.2 or later. Clients must also support cipher suites with perfect forward secrecy (PFS) such as Ephemeral Diffie-Hellman (DHE) or Elliptic Curve Ephemeral Diffie-Hellman (ECDHE). Most modern systems such as Java 7 and later support these modes.

All parties involved in the communication authenticate each other using TLS, IAM policies and tokens. The communication between the Agents and the replication server is based on TLS 1.2 only with the highest standard of cipher suite (PFS, ECDHE). Requests between the agent and AWS

Transform MGN as well as between the replication server and Application Migration Service are signed using an access key ID and a secret access key that is associated with an IAM principal.

All requests must be signed using the [AWS Security Token Service](#) (AWS STS), which allows you to generate temporary security credentials to sign requests. Alternatively, use credentials that are associated with an IAM principal.

AWS Transform MGN customers must ensure that they manually delete their access keys after installing the AWS Replication Agent and successful migration. AWS does not delete these keys automatically. AWS Transform MGN does delete the keys from source servers after they are disconnected from the service. If you want your keys to automatically stop working at a certain date after you have finished using them so that you do not have to worry about manually deleting them, you can do so through the [IAM permissions boundary](#) and the [aws:CurrentTime global context key](#).

AWS Transform MGN customers should use [Amazon EBS encryption](#).

AWS Transform MGN customers should secure their replication servers by reducing their exposure to the public internet. This can be done through:

1. Using security groups to only allow permitted IP addresses to connect to the replication servers. [Learn more about Security Groups](#).
2. Using a VPN to connect to the replication servers, such as the AWS site-to-site VPN. [Learn more about the AWS Site-to-site VPN](#).

AWS Transform MGN creates and uses the "aws-replication" user within the Linux Source server. The AWS Transform MGN replication server and AWS Replication Agent run under this user. Although this is not a root user, this user needs to be part of the disk group that grants this user full read and write permissions to block devices.

 Note

AWS Transform MGN only uses these permissions to read from block devices.

Compliance validation for AWS Transform MGN

Third-party auditors assess the security and compliance of AWS Transform MGN as part of multiple AWS compliance programs.

For a list of AWS services in scope of specific compliance programs, see [AWS Services in Scope by Compliance Program](#) . For general information, see [AWS Compliance Programs](#) .

You can download third-party audit reports using AWS Artifact. For more information, see [Downloading Reports in AWS Artifact](#) .

Your compliance responsibility when using AWS Transform MGN is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. AWS provides the following resources to help with compliance:

- [Security and Compliance Quick Start Guides](#) – These deployment guides discuss architectural considerations and provide steps for deploying security- and compliance-focused baseline environments on AWS.
- [Architecting for HIPAA Security and Compliance Whitepaper](#) – This whitepaper describes how companies can use AWS to create HIPAA-compliant applications.
- [AWS Compliance Resources](#) – This collection of workbooks and guides might apply to your industry and location.
- [Evaluating Resources with Rules](#) in the *AWS Config Developer Guide* – AWS Config; assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.
- [AWS Security Hub CSPM](#) – This AWS service provides a comprehensive view of your security state within AWS that helps you check your compliance with security industry standards and best practices.

Cross-service confused deputy prevention

The confused deputy problem is a security issue where an entity that doesn't have permission to perform an action can coerce a more-privileged entity to perform the action. In AWS, cross-service impersonation can result in the confused deputy problem. Cross-service impersonation can occur when one service (the *calling service*) calls another service (the *called service*). The calling service can be manipulated to use its permissions to act on another customer's resources in a way it should not otherwise have permission to access. To prevent this, AWS provides tools that help you protect

your data for all services with service principals that have been given access to resources in your account.

We recommend using the [aws:SourceArn](#) and [aws:SourceAccount](#) global condition context keys in resource policies to limit the permissions that AWS Transform MGN gives another service to the resource. If you use both global condition context keys, the `aws:SourceAccount` value and the account in the `aws:SourceArn` value must use the same account ID when used in the same policy statement.


The value of `aws:SourceArn` must be `"arn:aws:mgn:*:123456789012:source-server/*"`

The most effective way to protect against the confused deputy problem is to use the `aws:SourceArn` global condition context key with the full ARN of the resource. If you don't know the full ARN of the resource or if you are specifying multiple resources, use the `aws:SourceArn` global condition context key with wildcards (*) for the unknown portions of the ARN. For example, `arn:aws:servicename::123456789012:*` .

The following example shows how you can use the `aws:SourceArn` and `aws:SourceAccount` global condition context keys in AWS Transform MGN to prevent the confused deputy problem.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ConfusedDeputyPreventionExamplePolicy",
      "Effect": "Allow",
      "Principal": {
        "Service": "mgn.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringLike": {
          "aws:SourceAccount": "123456789012"
        },
        "ArnLike": {
          "aws:SourceArn": "arn:aws:mgn:*:123456789012:source-server/*"
        }
      }
    }
  ]
}
```



Troubleshooting

This section provides help for common communication, credential, installation, and replication issues with AWS Transform MGN.

Topics

- [Troubleshooting launch errors](#)
- [Troubleshooting communication errors](#)
- [Troubleshooting agent issues](#)
- [Troubleshooting agentless replication issues](#)
- [Common replication errors](#)
- [Other troubleshooting topics](#)
- [Target instance cannot connect to Active Directory after migration](#)
- [Troubleshooting FSx for ONTAP issues](#)

Troubleshooting launch errors

Use the information in this section to troubleshoot launch errors.

Topics

- [Do I need to recreate the entire launch template for every version?](#)
- [Error - AccessDeniedException - Must be admin user](#)
- [VPCIdNotSpecified error](#)
- [Error: Failed to connect using HTTP channel](#)
- [Could not take up-to-date snapshot. Launching from snapshot taken on...](#)

Do I need to recreate the entire launch template for every version?

When you save a new template version, it is tagged as the latest version. However, for a multitude of reasons, AWS Transform MGN uses the version marked as the default for its purposes. So to actually have MGN recognize the changes you make, you need to go into the template itself, and change the default version to the version you have just updated.

1. Create the new template version.
2. Select the launch template in the success message, and then select **Actions** and choose **Set default version**.
3. From the drop down menu select the latest version, and then choose **Set as default version**.

Error - AccessDeniedException - Must be admin user

If you receive an AccessDeniedException error when attempting to log into AWS Transform MGN for the first time and set up your replication settings template, it means that you are not the administrator of your AWS Account and therefore cannot initialize MGN. You must be the Admin user of your AWS Account to initialize MGN. [Learn more about initializing MGN.](#)

VPCIdNotSpecified error

The EC2 launch template does not automatically set a specific Subnet. As such, EC2 will attempt to launch in a Subnet within the default VPC. If you have removed your default VPC, EC2 will fail to launch any instance for which there is no valid Subnet specified. Ensure that you specify a subnet if that is the case, or AWS Transform MGN instance launch will fail. You may see the VPCIdNotSpecified error if:

- A default subnet/VPC is not selected in the EC2 launch template.
- An incorrect target subnet is specified in the EC2 launch template.
- the EC2 launch template with the correct subnet settings is not set as the default.

Error: Failed to connect using HTTP channel

This error mostly occurs when the conversion server is unable to communicate with the necessary AWS Endpoints for [staging area communication](#).

- Check if any network changes were made in the staging area that could affect the conversion server reaching the AWS Endpoints (firewall settings, DNS settings, security group settings, route table settings, and access control list settings).
- Test TCP Port 443 connectivity with a test instance from the staging area subnet, to the [required endpoints](#).
- If the issue persists after confirming network connectivity, [create a case](#) with AWS Premium Support for further investigation.

Could not take up-to-date snapshot. Launching from snapshot taken on...

When a test or cutover instance is launched, AWS Transform MGN will attempt to create the latest consistent snapshot of the source server. MGN will wait for all the snapshots to become available and once they are ready, will proceed with the launch workflow.

If you see a timeout message when launching a test or cutover instance, it means the snapshot creation timed out. In this case, MGN will use the latest successful snapshot for that source server to launch the instance. This ensures you can still launch an instance, but the instance will only contain data current up to the timestamp specified in the message.

To launch a test or cutover instance with the most up-to-date data, determine why the latest snapshot could not be created. Common causes include the source server not having a "Healthy" status, or backlog/lag.

Also check the CloudTrail Event History for errors on the CreateSnapshot and DescribeSnapshot API calls, which can prevent timely EBS snapshot creation. Resolving these underlying issues will allow successful creation of up-to-date snapshots for test and cutover instances.

Troubleshooting communication errors

Use the information in this section to troubleshoot communication errors.

Topics

- [Solving communication problems over TCP Port 443 between the staging area and the AWS Transform MGN](#)
- [Authenticate with service errors](#)
- [Calculating the required bandwidth for TCP Port 1500](#)
- [Verifying communication over Port 1500](#)
- [Solving communication problems over Port 1500](#)
- [Missing Amazon Linux 2023 repository access](#)

Solving communication problems over TCP Port 443 between the staging area and the AWS Transform MGN

- **DHCP** – [Check the DHCP options set of the VPC of the staging area.](#)

Ensure that the IPv4 CIDR, the DHCP options set, the route table, and the network ACL are correct.

- **DNS** – Ensure that your security groups allow outbound DNS resolution over TCP Port 53, and outbound HTTPS connectivity over TCP Port 443.
- **Route Rules** – the route rules on the staging area subnet may be inaccurately set. The route rules should allow outbound traffic to the Internet.

To check and set the route rules on the staging area subnet:

1. Sign in to [AWS console](#), choose **Services** and select **VPC** under **Networking & Content Delivery**.
2. On the **VPC Dashboard** toolbar, select the **Route Tables** option.
3. On **Route Tables** page, check the box of the route table of your staging area.
4. This will open the details for your Route Table. Navigate to the **Routes** tab.
5. Within the **Target** column of the **Routes** tab, find the route you are using for the outbound communication to the Internet (either **igw**- Internet Gateway, **vgn** - **VPN** or **i** - EC2 instance). Verify that the address space in the Destination column is allowing access to Amazon S3, Amazon EC2, and MGN in the AWS Region.
6. If the address is not **0.0.0.0/0**, change it to **0.0.0.0/0**.

Choose the **Edit** button.

7. Input **0.0.0.0/0** into the Destination field for the correct **Target**. Choose **Save**.

Note: If you are using VPN, enter a specific IP address range in the **Destination** column.

- **Network ACL** – The network ACL on the staging area subnet may block the traffic. Verify that the ephemeral ports are open.

Authenticate with service errors

The replication server needs to be able to reach the MGN endpoint and have the proper IAM permissions. This can fail for a number of reasons, including:

- The replication server is in a subnet without access to VPC endpoints for MGN or the [public endpoints](#).
- In some use cases, when using a custom DNS server, DNS traffic shifts to TCP instead of UDP. The solution for this is to [update](#) the [Migration Service Security Group](#) to allow outbound TCP traffic on port 53.
- The Replication Server does not have the proper [IAM policy](#).

Calculating the required bandwidth for TCP Port 1500

The required bandwidth for transferring the replicated data over TCP Port 1500 should be based on the write speed of the participating source servers. The recommended bandwidth should be at least the sum of the average write speed of all replicated source servers.

Minimal bandwidth = the sum of the average write speeds of all Source servers

For example, suppose you are replicating two source servers. One has a write speed of 5 MBps (meaning 5 megabytes of data every second), while the other has 7 MBps. In this case, the recommended bandwidth should be at least 12 MBps.

Finding the write speed of your source servers

To calculate the required bandwidth for transferring replicated data over TCP Port 1500, you need to know the write speed of your source servers. Use the following tools to find the write speed of your source servers:

Linux

Use the `iostat` command-line utility, located in the `systat` package. The `iostat` utility monitors system input/output device loading and generates statistical reports.

The links below lead to third-party websites not affiliated with or endorsed by AWS. The content on these external sites has not been reviewed or verified by AWS.

The `iostat` utility is installed [with yum](#) (RHEL/CentOS), via [apt-get](#) (Ubuntu) and via [zypper](#) (SUSE.)

To use `iostat` for checking the write speed of a Source Server, enter the following:`iostat -x <interval>`

- `-x` - displays extended statistics.
- `<interval>` - the number of seconds `iostat` waits between each report. Each subsequent report covers the time since the previous report.

For example, to check the write speed of a machine every 3 seconds, enter the following command:

```
iostat -x 3
```

We recommend that you run the `iostat` utility for at least 24 hours, since the write speed to the disk changes during the day, and it will take 24 hours of runtime to identify the average running speed.

Windows

Choose one of these options to determine the read/write speed:

- Open Resource Monitor `perfmon.exe /res` and view the disk activity on the **Disk** tab.
- Open Performance Monitor `perfmon.exe` and add the **Avg. Disk Bytes/Write** or **Avg. Disk Bytes/Read** counter for Logical Disk or Physical Disk.

Verifying communication over Port 1500

If there is a connection problem from the Source server to the Replication Servers or the Staging Area, use the following methods to check the connection.

To verify the integrity of the connection from a Source server to the Staging Area over TCP Port 1500:

1. Launch a new Linux machine in the Staging Area subnet.
2. On the new Linux machine, run the following command to open a listener in the Staging Area subnet:

```
nc -l 1500
```

3. On the Source Server, run the following command to check connectivity:

```
telnet <new machine ip> 1500
```

Solving communication problems over Port 1500

To solve connectivity problems between Source server and the staging area, check the following:

- The Network ACL on the staging area subnet may deny the traffic.
- Route Rules on the staging area subnet may be inaccurately set.
- The firewall, both internal and external, in the Source Server/infrastructure may block communication.
- The **Use VPN...** checkbox in AWS Transform MGN console may not be set correctly.

Enabling the network ACL

The Network ACL on the staging area subnet may block connectivity. By default, the network ACL allows connectivity. However, if the ACL setting was changed to deny traffic, you need to change it back.

To check and activate the network ACL on the staging area subnet:

1. Sign in to the AWS console, choose **Services** and select **VPC** under **Networking & Content Delivery**.
2. On the **Resources** list, select the **Network ACL** option:
3. On **Network ACL** page, select the check box next to the Network ACL of your staging area.
4. On the details table of the selected **Network ACL**, select the **Inbound Rules** tab.
5. On the **Inbound Rules** tab, verify that the rule that determines the traffic to replication server subnet set to **Allow**.

Note: The target should allow traffic on TCP Port 1500 from the address space of the source environment. The Network ACL does not necessarily need to be open to all port ranges, as in the screenshot below.

6. If the rule is set to **Deny**, choose **Edit**.
7. Choose the dropdown under **Allow/Deny** and select **Allow**. Choose **Save**.
8. You will also need to check the **Ephemeral Ports** on the **Outbound Rules** tab. Within the same **Network ACL**, navigate to the **Outbound Rules** tab.
9. You will need to ensure that you are allowing the correct **Ephemeral Port range** for your particular client. [Ephemeral Port range varies based on each client's operating system](#). Choose the Edit button to edit your **Ephemeral Port's Port Range** category.

10. Edit the **Port Range** and choose **Save**. You may have to create a new Rule by choosing the **Add another rule** button.

Setting route rules on the staging area subnet

To check and set the route rules on the staging area subnet in AWS:

1. Sign in to AWS console, choose **Services** and select **VPC** under **Networking & Content Delivery**.
2. On the **VPC Dashboard** toolbar, select the **Route Tables** option.
3. On the **Route Tables** page, check the box of the Route Table of your staging network.
4. This will open the details for your Route Table. Navigate to the **Routes** tab.
5. Within the **Target** column of the **Routes** tab, find the route you are using for the inbound traffic from the **Source** on TCP Port 1500 (either **igw** - Internet Gateway, **vgnw** - VPN or **i** – EC2 instance). Verify that the **Destination address** is **0.0.0.0/0**.

Note: The Rule may be specific to the address space of the source servers.

6. If the address is not 0.0.0.0/0, you will need change it to 0.0.0.0/0.


Note: The Rule may be specific to the address space of the Source Servers.

1. Choose the Edit button.
2. Input **0.0.0.0/0** into the **Destination** field for the correct target. Choose **Save**.

Note: If you are using VPN, enter a specific IP address range in the **Destination** column.

Firewall (both internal and external) in the source server/infrastructure.

Firewall issues may have several causes. Check the following if you experience any firewall issues, such as Windows Firewall connection issues:



- Ensure that the subnet you assigned for the replication servers still exists.

Missing Amazon Linux 2023 repository access

Replication servers and conversion servers are based on Amazon Linux 2023 (AL2023) and require access to the AL2023 package repository hosted in Amazon S3. If the staging area subnet cannot reach this repository, replication or conversion may fail.

Symptoms of this issue include:

- Replication stalling or failing after the replication server is launched.
- Replication server unable to install or update required packages.
- Conversion server failing during the boot conversion process.

Cause: Amazon S3 VPC gateway endpoint policy missing AL2023 bucket (isolated subnets)

If the staging area subnet does not have outbound internet access and uses an Amazon S3 VPC gateway endpoint, the endpoint policy must allow access to the AL2023 repository bucket. Add the following resource ARN to the endpoint policy:

```
"arn:aws:s3:::al2023-repos-<REGION>-de612dc2/*"
```

Cause: Firewall or proxy blocking the AL2023 repository URL (firewall-restricted environments)

If the staging area subnet has outbound internet access through a firewall or proxy, ensure that the AL2023 repository URL is allowlisted. The AL2023 repository is accessed through the Amazon S3 dual-stack endpoint:

```
https://al2023-repos-<region>-de612dc2.s3.dualstack.<region>.amazonaws.com/
```

For the complete list of Amazon S3 buckets required by AWS Transform MGN and the example Amazon S3 VPC endpoint policy, see [Communication between the staging area subnet and S3](#).

Troubleshooting agent issues

Use the information in this section to troubleshoot issues with installing the replication agent.

Topics

- [Error: Installation failed](#)
- [Where can I find the MGN Agent logs?](#)

Error: Installation failed

This type of error means that the agent was not installed on the source server, and therefore the server will not appear on the AWS Transform MGN console. After you fix the issue that caused the installation to fail, you need to rerun the Agent Installer file to install the agent.

Error: too many volumes selected for replication

Installation fails with an error similar to the following when you select more disks than the AWS Replication Agent supports:

```
The devices you entered span over more than 63 volumes. Up to 63 volumes are supported  
per machine.  
Installation failed.
```

The AWS Replication Agent can replicate up to 63 disks from a single source server. To resolve this error, select 63 or fewer disks for replication and rerun the Agent Installer file.

This app can't run on your PC error - Windows

If you encounter the following error "This app can't run on your PC", when trying to install the AWS Replication Agent on your Windows 10 source Server, try the following.

This error is indicative that your particular version of Windows 10 is likely the 32-bit version. To verify this, you can

1. Use the Windows key + I keyboard shortcut to open the Settings app.
2. Choose System.
3. Choose About.
4. Under System type, you will see two pieces of information: if it says 32-bit operating system, x64-based processor, then it means that your PC is running a 32-bit version of Windows 10 on a 64-bit processor.

If it says 32-bit operating system, x86-based processor, then your computer doesn't support Windows 10 (64-bit).

If your OS is indeed 64-bit, then there may be other elements blocking the installation of your agent. The block is actually coming from the Windows Operating System itself. You would need to identify what the cause is. One of the ways is running [sfc scan](#).

Is having a mounted '/tmp' directory a requirement for the agent?

The simple requirement is just to have enough free space. There is no need for this to be a separate mount. The need for the '/tmp' requirement is actually only if '/tmp' is a separate mount. If '/tmp' is not a separate mount, then it would fall under '/', for which we have the 2 GB free requirement. This allows for the '/tmp' to fall into this requirement.

Installation failed - old agent

Installation may fail due to an old AWS Replication Agent. Ensure that you are attempting to install the latest version of the AWS Replication Agent. You can learn how to download the Agent [here](#).

Installation cannot be completed - CloudEndure Agent

Agent installation will fail if the source server already has the CloudEndure User Agent installed on it. You will need to [uninstall the CloudEndure Agent](#) and then install the AWS Replication Agent in order to proceed.

At times, uninstalling the CloudEndure Agent alone is not enough, as the Agent driver may remain. If that is the case, you will need to delete the agent driver manually.

Linux

Run the following command to identify the CloudEndure driver:

```
lsmod | grep CE_AgentDriver
```

Then, run the following command to delete the driver if it exists:

```
rmmod CE_AgentDriver
```

Windows

Run the following command in cmd to identify the CloudEndure driver:

```
sc query ce_driver
```

```
sc query ce_filter_driver
```

Then, run the following command to delete the driver if it exists:

```
sc delete ce_driver
```

```
sc delete ce_filter_driver
```

Installation failed on Linux Server

If the installation failed on a Linux Source server, check the following:

1. Free Disk Space

Free disk space on the root directory – verify that you have at least 2 GB of free disk on the root directory (/) of your Source Server. To check the available disk space on the root directory, run the following command: `df -h /`

Free disk space on the /tmp directory – for the duration of the installation process only, verify that you have at least 1 GB of free disk on the /tmp directory. To check the available disk space on the /tmp directory run the following command: `df -h /tmp`

After you have entered the above commands for checking the available disk space, the results will be displayed as follows:


```
ubuntu@Linux-1:~$ df -h /
Filesystem      Size  Used Avail Use% Mounted on
/dev/xvda1      7.8G  1.4G  6.0G  19% /
ubuntu@Linux-1:~$ df -h /tmp
Filesystem      Size  Used Avail Use% Mounted on
/dev/xvda1      7.8G  1.4G  6.0G  19% /tmp
```

2. The format of the list of disks to replicate

During the installation, when you are asked to enter the disks you want to replicate, do NOT use apostrophes, brackets, or disk paths that do not exist. Type only existing disk paths, and separate them with a comma, as follows:

```
/dev/xvda,/dev/xvdb.
```

3. Version of the Kernel headers package

Verify that you have kernel-devel/linux-headers installed that are exactly of the same version as the kernel you are running.

The version number of the kernel headers should be completely identical to the version number of the kernel. To handle this issue, follow these steps:

a. Identify the version of your running kernel.

To identify the version of your running kernel, run the following command:

```
uname -r
```

```
[root@ip-172-31-1-164 ~]# uname -r
4.4.41-36.55.amzn1.x86_64
[root@ip-172-31-1-164 ~]#
```

The 'uname -r' output version should match the version of one of the installed kernel headers packages (kernel-devel-<version number> / linux-headers-<version number>).

b. Identify the version of your kernel-devel/linux-headers.

To identify the version of your running kernel, run the following command:

On RHEL/CENTOS/Oracle/SUSE:

```
rpm -qa | grep kernel
```

```
[root@ip-172-31-1-164 ~]# rpm -qa | grep kernel
kernel-tools-4.4.41-36.55.amzn1.x86_64
kernel-4.4.41-36.55.amzn1.x86_64
kernel-headers-4.4.41-36.55.amzn1.x86_64
[root@ip-172-31-1-164 ~]#
```

Note: This command looks for kernel-devel.

On Debian/Ubuntu: apt-cache search linux-headers

```
ubuntu@Linux-1:~$ apt-cache search linux-headers
linux-headers-3.13.0-24 - Header files related to Linux kernel version
3.13.0
linux-headers-3.13.0-24-generic - Linux kernel headers for version 3.1
3.0 on 64 bit x86 SMP
linux-headers-3.13.0-24-lowlatency - Linux kernel headers for version
3.13.0 on 64 bit x86 SMP
```

c. Verifying that the folder that contains the kernel-devel/linux-headers is not a symbolic link.

Sometimes, the content of the kernel-devel/linux-headers, which match the version of the kernel, is actually a symbolic link. In this case, you will need to remove the link before installing the required package.

To verify that the folder that contains the kernel-devel/linux-headers is not a symbolic link, run the following command:

On RHEL/CENTOS/Oracle/SUSE:

```
ls -l /usr/src/kernels
```

On Debian/Ubuntu:

```
ls -l /usr/src
```

```
ubuntu@Linux-1:~$ ls -l /usr/src
total 8
lrwxrwxrwx 1 root root 41 May 29 15:40 3.13.0-116-generic -> /usr/src/linux-headers-3.13.0-116-generic
drwxr-xr-x 24 root root 4096 Apr 5 20:43 linux-headers-3.13.0-116
drwxr-xr-x 7 root root 4096 Apr 5 20:43 linux-headers-3.13.0-116-generic
ubuntu@Linux-1:~$
```

In the above example, the results show that the linux-headers are not a symbolic link.

d. [If a symbolic link exists] Delete the symbolic link.

If you found that the content of the `kernel-devel/linux-headers`, which match the version of the kernel, is actually a symbolic link, you need to delete the link. Run the following command:

```
rm /usr/src/<LINK NAME>
```

For example: `rm /usr/src/linux-headers-4.4.1`

e. Install the correct `kernel-devel/linux-headers` from the repositories.

If none of the already installed `kernel-devel/linux-headers` packages match your running kernel version, you need to install the matching package.

Note: You can have several kernel headers versions simultaneously on your OS, and you can therefore safely install new kernel headers packages in addition to your existing ones (without uninstalling the other versions of the package.) A new kernel headers package does not impact the kernel, and does not overwrite older versions of the kernel headers.

Note: For everything to work, you need to install a kernel headers package with the exact same version number of the running kernel.

To install the correct `kernel-devel/linux-headers`, run the following command:

On RHEL/CENTOS/Oracle/SUSE:

```
sudo yum install kernel-devel-`uname -r`
```

On Oracle with Unbreakable Enterprise Kernel:

```
sudo yum install kernel-uek-devel-`uname -r`
```

On Debian/Ubuntu:

```
sudo apt-get install linux-headers-`uname -r`
```

f. [If no matching package was found] Download the matching `kernel-devel/linux-headers` package.

If no matching package was found on the repositories configured on your server, you can download it manually from the Internet and then install it.

To download the matching *kernel-devel/linux-headers* package, navigate to these sites:

- [RHEL and Centos](#)
- [Oracle Linux](#)
- [SUSE](#)
- [Debian](#)
- [Ubuntu](#)

4. The make, openssl, wget, curl, gcc and build-essential packages.

Note: Usually, the existence of these packages is not required for Agent installation. However, in some cases where the installation fails, installing these packages will solve the problem.

If the installation failed, the make, openssl, wget, curl, gcc, and build-essential packages should be installed and stored in your current path.

To verify the existence and location of the required packages, run the following command:

which <package>

For Example, to locate the make package:

which make

```
[root@ip-172-31-1-164 ~]# which make
/usr/bin/make
[root@ip-172-31-1-164 ~]#
```

5. Error: urlopen error [Errno 10060] Connection times out.

This error occurs when outbound traffic is not allowed over TCP Port 443. Port 443 needs to be open outbound to the MGN Service endpoint.

6. Powerpath support

Contact AWS Support for instructions on how to install the AWS Transform MGN Agent on such machines.

7. Error: You need to have root privileges to run this script.

Make sure you run the installer either as root or by adding sudo at the beginning:

```
sudo ./aws-replication-installer-init
```

8. Error: version `GLIBC_2.7' not found (required by ./aws-replication-installer-64bit)

You receive this error when you try to install the agent on an unsupported Linux operating system. See [Supported Linux operating systems](#).

Installation failed on Windows machine

If the installation failed on a Windows Source server, check the following:

1. .NET Framework

Verify that .NET Framework version 3.5 or above is installed on your Windows Source servers.

2. Free disk space

Verify that there is at least 1 GB of free disk space on the root directory (C:\) of your Source servers for the installation.

3. net.exe and sc.exe location

Verify that the net.exe and/or sc.exe files, located by default in the C:\Windows\System32 folder, are included in the **PATH Environment Variable**.

- a. Navigate to **Control Panel >System and Security >System >Advanced system settings**.
- b. On the **System Properties** dialog box **Advanced** tab, choose the **Environment Variables** button.
- c. On the **System Variables** section of the **Environment Variables** pane, select the **Path** variable. Then, choose the **Edit** button to view its contents.
- d. On the **Edit System Variable** pane, review the defined paths in the **Variable value** field. If the path of the net.exe and/or sc.exe files does not appear there, manually add it to the **Variable value** field, and choose **OK**.

Windows - Installation Failed - Request Signature

If the AWS Replication Agent installation fails on Windows with the following error:

```
botocore.exceptions.ClientError: An error occurred (InvalidSignatureException) when calling the GetAgentInstallationAssetsForMgn operation: {"message": "The request signature we calculated does not match the signature you provided. Check your AWS Secret Access Key and signing method. Consult the service documentation for details."}
```

Attempt to rerun the installer with PowerShell instead of CMD. At times, when the installer is run in CMD, the AWS Secret Key does not get pasted properly into the installer and causes installation to fail.

Error – certificate verify failed

This error (CERTIFICATE_VERIFY_FAILED) may indicate that the OS does not trust the certification authority used by our endpoints. To resolve this issue, try the following steps:

1. Open Microsoft Edge or Internet Explorer to update the operating system trusted root certificates. This will work if the operating system does not have restrictions to download the certificates.
2. If the first step does not resolve the issue, [download and install the Amazon Root Certificates manually](#).

Where can I find the MGN Agent logs?

The MGN Agent logs are stored in agent.log.0:

- **Linux:** /var/lib/aws-replication-agent/agent.log.0
- **Windows 64 bit:** C:\Program Files (x86)\AWS Replication Agent\agent.log.0
- **Windows 32 bit:** C:\Program Files\AWS Replication Agent\agent.log.0

In addition, you can review the installation log located in: <install_path>\aws_replication_agent_installer.log

Troubleshooting agentless replication issues

Discovery related troubleshooting

- No machines are discovered -

Verify vCenter client is still connected to the service and able to communicate with the required endpoints.

Verify vCenter user permissions, which you have configured the MGN Agentless client with, are correct by reviewing the [agentless snapshot based replication section](#).

- The same Machine is discovered twice -

Source Servers are identified by their vCenter UUID. If a server changes its UUID (which may happen in edge cases), the same guest can appear as two source servers, where only one of them is responsive (the new UUID). In this case, you may archive the previous source server, or use the API to delete it.

- Two machines with the same UUID -

Both replications are stalled, duplicate UUID reported to the service log. The workaround is to change the UUID of one of the machines. [Learn more about changing UUID](#)

Agentless Replication related troubleshooting

- Replication state is STALLED -

Verify vCenter client is still connected to the service and able to communicate with the required endpoints.

- dataReplicationError is UNSUPPORTED_VM_CONFIGURATION -

Try to delete all VMWare snapshots on the guest, and wait for the next iteration. This usually works around snapshotting and CBT issues

- dataReplicationError is LAST_SNAPSHOT_JOB_FAILED

Check to see if the snapshots are successfully being created in vCenter. You can check this by choosing the VM which is in the error state > choosing the Monitor tab > Tasks

Check the replication log which can be found at `/var/lib/aws-vcenter-client/active/tmp/Unique-ID/snapshot-num.log.0`

Logs

`/var/lib/aws-vcenter-client/active/aws-vcenter-client-upgrader.log` - This is the log for the upgrade process of the Agentless client software

`/var/lib/aws-vcenter-client/active/vcenter_commands_client.log` - This is the main agentless client log that will show the log of the agentless discovery process, etc.

`/installation-directory/aws-vcenter-client-installer.log` - This log will show any issues with the installation process of the MGN agentless client. This will be located in the directory that you have executed the installer from.

/var/lib/aws-vcenter-client/active/tmp/*Unique-ID/snapshot-num.log.0* - When you start replication for a discovered server, it will create a snapshot in vCenter and generate this log which is similar to the agent log. For each snapshot, there is a different log. This will show a log of replication process for the snapshot.

Common replication errors

This section describes common replication errors, possible explanations, and potential mitigations.

Replication errors

- [Agent not seen](#)
- [Not converging](#)
- [Snapshot failure](#)
- [Unstable network](#)
- [Failed to connect AWS replication Agent to replication software](#)
- [Failed to establish communication with replication software](#)
- [Failed to create firewall rules](#)
- [Failed to create security groups](#)
- [Failed to authenticate with service](#)
- [Failed to create staging disks](#)
- [Replication stuck at Attach Staging Disks](#)
- [Failed to pair the replication agent with replication server](#)
- [Stalled replication when replicating a source volume smaller than 1MiB](#)
- [Unknown data replication error](#)

Agent not seen

- If you see this message, ensure that:
 - The source Server has access to the AWS Transform MGN service.
 - The replication agent is in running state. For Windows, use Windows services management console (services.msc) or command line (for example, get-services PowerShell). For Linux, use the command `systemctl status aws-replication`.

If the agent is indeed in running state, verify that the connectivity to the Regional MGN endpoint on TCP Port 443. [Learn more about verifying connectivity to MGN regional endpoints.](#)

Not converging

This error message (NOT_CONVERGING) could indicate an inadequate replication speed.

Not converging error implies that there is a backlog, but the transfer of data in comparison to the growth of the data on the source server is slower. If the source server is writing more to the disk as compared to the speed at which its sending the data then we get the not converging error.

- Follow the instructions on [calculating the required bandwidth](#).
- [Verify network bandwidth](#).
- Verify replicator Amazon EBS volumes (associated with the source server) performance. If required, modify EBS volume type from the MGN console: Go to the specific source server page and select the **Disk settings** tab.
- Verify the source server performance. For example CPU and Memory utilization.

Snapshot failure

This error message (SNAPSHOTS_FAILURE) indicates that the service is unable to take a consistent snapshot.

This can be caused by:

- Inadequate IAM permissions – Check your CloudTrail logs for any errors in the CreateSnapshot API call. Ensure that you have the required IAM permissions (attached to the required IAM roles).

Restrictive Service Control Policies – Check if your AWS Organization has a Service Control Policy (SCP) that is preventing the snapshot creation.
- API throttling – Check your CloudTrail logs for API throttling errors.

Unstable network

This error message (UNSTABLE_NETWORK) may indicate that there are network issues. Check your connectivity, then [run the network bandwidth test](#).

Failed to connect AWS replication Agent to replication software

This error message (FAILED_TO_PAIR_AGENT_WITH_REPLICATION_SOFTWARE) may indicate a pairing issue. AWS MGN needs to provide the replication server and agent with information to allow them to communicate. Make sure there is network connectivity between the agent, migration server, and the MGN endpoint.

If the issue persists, contact support.

Failed to establish communication with replication software

This error message (FAILED_TO_ESTABLISH_AGENT_REPLICATOR_SOFTWARE_COMMUNICATION) may suggest that there are network connectivity issues. Make sure you have network connectivity between the agent, replication server and the MGN endpoint.

Failed to create firewall rules

This error message (Firewall rules creation failed) can be caused by several reasons.

1. Ensure that the IAM permission prerequisites are met.
2. Review the replication settings of the associated source server.

Failed to create security groups

AWS MGN could not create or update the security group for the replication server. This is usually one of the following:

- A Service Control Policy (SCP) or permissions boundary in your AWS Organization is blocking the MGN service-linked role from modifying security groups.
- A security group named **AWS Application Migration Service default Replication Server Security Group** already exists in the staging area VPC but was not created by MGN. Remove or rename it so MGN can create and manage its own.

Failed to authenticate with service

This error message (Failed to authenticate the replication server with the service) may indicate a communication issue between the replication server and the MGN endpoint on TCP Port 443. Check the subnet you selected and ensure that TCP Port 443 is open from your replication server.

To verify the connection:

- Launch a test Amazon Linux 2 EC2 instance in the same subnet that was selected in the replication settings.
- On the server, run the following command:

```
wget <enter_MGN_regional_endpoint>
```

- If the command fails, there is a connectivity problem.

Failed to create staging disks

This error message (Failed to create staging disks) may indicate that your AWS account is configured to use encrypted EBS disks, but the IAM user does not have the required permissions to encrypt using the selected KMS key.

Check your CloudTrail logs for any errors in the CreateVolume API call. Then ensure that you have the required IAM permissions attached to the specified IAM role. If the issue persists, also check your KMS Key Policy for any statements that may prevent MGN from using the selected KMS key.

Replication stuck at Attach Staging Disks

If replication is stuck at the "Attach Staging Disks" stage during initial sync, the replication server may have reached the maximum number of Amazon EBS volumes that can be attached to the instance type.

Cause: By default, MGN uses a single replication server per 15 source disks. If the total number of disks across source servers sharing a replication server exceeds the Amazon EBS volume attachment limit of the replication server instance type, the attach operation will stall.

Resolution:

- Use a replication server instance type that supports more Amazon EBS volume attachments. For a list of instance types and their supported volume counts, see [General purpose instances – EBS specifications](#).
- Use [dedicated replication servers](#) to reduce the number of disks per replication server.
- Reduce the number of disks being replicated per source server by excluding unnecessary volumes during agent installation.

Failed to pair the replication agent with replication server

This error message (Failed to pair replication agent with replication server) may be caused by multiple reasons. Make sure that you have connectivity between the replication agent, the replication server, and the MGN endpoint. If the issue persists, contact Support.

Stalled replication when replicating a source volume smaller than 1MiB

Replication is stalled which point to a common network, however in this edge case it indicates you are trying to replicate a disk on source server which is smaller than 1MiB. This is not supported. To mitigate, a. [Reinstall](#) the agent with manually identifying disks to replication, and not include the disk with size under 1MiB. b. Expand the volume and reinstall the agent.

Unknown data replication error

Unknown errors (unknown_error) can occur for any number of reasons. There are several steps you can take to attempt to mitigate the issue:

- Check connectivity.
- Check throttling.
- Check performance issue on the replication server.
- Check the [network bandwidth](#) between the agent and the replication server.
- [Check the replication agent logs.](#)

Other troubleshooting topics

Use the information in this section to help you with other troubleshooting.

Topics

- [Re-initialize the AWS Transform MGN](#)
- [Windows license activation – AWS](#)
- [Migration leaving behind replication volumes after cutover](#)
- [Replication lag issues](#)
- [Windows drive letter changes](#)

- [Windows Dynamic Disk troubleshooting](#)
- [Deleting MGN resources](#)
- [Set UEFI boot mode](#)

Re-initialize the AWS Transform MGN

AWS Transform MGN can be re-initialized in case of any issues with IAM service roles

To re-initialize the MGN service, follow these steps:

- Open the AWS Transform MGN Console and navigate to the correct region you are migrating to.
- In the left navigation pane, select "Settings". Under "Replication template," choose "Reinitialize service permissions" and then choose "Confirm."

Windows license activation – AWS

AWS Transform MGN converts the Windows OS licenses to AWS Windows licenses and activates them against the AWS KMS.

If license activation failed, follow [this AWS guide](#) to resolve.

Migration leaving behind replication volumes after cutover

If you are seeing left behind replication volumes in AWS after running the cutover process, then ensure that the names of the replication volumes match those given to them by AWS Transform MGN.

Most likely, you have a script running in your AWS account that renames Amazon EBS volumes to match the name of the EC2 instance they are attached to.

However, by renaming an EBS volume used by MGN for replication you are severing its association with AWS MGN, and AWS MGN will not automatically clean up such volumes.

This can also occur if MGN service tags associated with EBS volumes are modified.

Key: AWSApplicationMigrationManaged; Value: mgn.amazonaws.com

Key: Name;Value: Application Migration Service Replication Volume

Replication lag issues

Potential solutions:

- Make sure that the Source server is up and running.
- Make sure that AWS MGN services are up and running.
- Make sure that TCP Port 1500 is not blocked outbound from the source server to the replication server.
- If the MAC address of the Source had changed, that would require a reinstallation of the AWS Replication Agent.
- If the source server was rebooted recently or the AWS Transform MGN was restarted, the disks are re-read after this and until it's finished, the Lag will grow.
- If the source server had a spike of write operations, the lag will grow until AWS Transform MGN manages to flush all the written data to the test or cutover instance replication server.
- Make sure you have selected the right replication instance and EBS type by using the following runbook: [AWSSupport-CalculateEBSPerformanceMetrics automation runbook](#) . The replication instance ID, which is used as an input parameter for the runbook, is available under the source server replication dashboard.

To learn more about replication lag troubleshooting, refer to this [AWS Support Knowledge Center article](#).

Windows drive letter changes

Users may see changes in Windows drive letter assignments (for example, Drive D changed to E) on target machines launched by AWS Transform MGN.

This happens because Windows sometimes re-configures the drive letters when a machine comes up on a new infrastructure, for example, if the Source server had a drive letter mapped to a disk that was not replicated (such as a network drive). You can solve this issue by remapping the drive letters on the test or cutover instance correctly after it has been launched.

Windows Dynamic Disk troubleshooting

Moving a Windows Dynamic Disk from a local computer to another computer may change the disk status to "Foreign", resulting in a disruption in replication. The solution is to import the foreign disk, as discussed in [this Microsoft troubleshooting article](#).

Deleting MGN resources

You can delete various MGN resources, including source servers, jobs, and the replication settings template, through the MGN API. Use the following API commands to delete resources:

DeleteSourceServer

Use the *DeleteSourceServer* API command to delete source servers.

This command:

- Deletes a single source server by ID.
- Successful deletion should result in a 204 HTTP response code.
- To delete source server the server should be in a *DISCONNECTED* or *CUTOVER* state. If the source server has not been cut over, you must first call the *DisconnectFromService* API and then call the *DeleteSourceServer* API.

DeleteJob

Use the *DeleteJob* API command to delete a Job.

This command:

- Deletes a single Job by ID.
- Successful deletion should result in a 204 HTTP response code.
- Job must be in a *COMPLETED* state to be deleted.

DeleteReplicationConfigurationTemplate

Use the *DeleteReplicationConfigurationTemplate* API command to delete the replication template.

This command:

- Deletes a single replication template by ID.
- Successful deletion should result in a 204 HTTP response code.
- All source servers and jobs must first be deleted before calling the *DeleteReplicationConfigurationTemplate* API.

Set UEFI boot mode

UEFI boot mode can only be used with Nitro based EC2 Instance types. Nitro requires ENA driver, which is only supported from kernel 3.10.

UEFI to UEFI boot mode is not supported with agentless replication.

If you get the error `Source server boot mode is UEFI which is inconsistent with target instance`. It might be because

- The OS uses an old UEFI format from kernel 3.8 or earlier. If so, set the source server boot mode to 'Legacy BIOS'.
- You are performing UEFI to UEFI boot mode with agentless replication.

Target instance cannot connect to Active Directory after migration

When you migrate domain-joined Windows servers, the target instance may fail to authenticate with Active Directory. This prevents login with domain credentials or access to domain resources.

Causes:

- **Network connectivity** – No network path between the target VPC and your AD domain controllers. This requires an [AWS Site-to-Site VPN](#) or [AWS Direct Connect](#), and security groups/ACLs must allow AD ports (TCP/UDP 389, 636, 88, 53, 445, 135, 3268, 3269).
- **DNS resolution** – MGN resets network settings to DHCP during conversion. The VPC's default AmazonProvidedDNS cannot resolve on-premises AD domain names, so the instance cannot locate domain controllers.

Resolution:

1. **Ensure network connectivity** – Verify a network path exists between the target VPC and your AD domain controllers. Confirm that [security groups](#), [network ACLs](#), and on-premises firewalls allow AD traffic.
2. **Configure DNS resolution** – Use one of these approaches:

- *Recommended:* Create a [Route 53 Resolver](#) outbound endpoint with a forwarding rule for your AD domain. This preserves AWS service endpoint resolution. See [Integrating DNS with Route 53 Resolvers](#).
- *Alternative:* Create a [custom DHCP options set](#) with your AD DNS servers. Note: this may break AWS service endpoint resolution unless your DNS servers also forward AWS domain queries.

Test by launching a test instance in the target VPC before performing a cutover migration.

Troubleshooting FSx for ONTAP issues

This section covers common issues when using FSx for ONTAP as the target storage type with MGN.

Topics

- [the section called “Troubleshooting FSx for ONTAP iSCSI connectivity”](#)
- [the section called “FSx for ONTAP replication errors”](#)
- [the section called “Failed to start data transfer”](#)
- [the section called “Not converging”](#)
- [the section called “Replication volume not deleted after Finalize cutover/Disconnect from service \(FlexClone split blocked by backup\)”](#)
- [the section called “Orphaned FSx for ONTAP target volumes \(FlexClone\) after launch cleanup”](#)
- [the section called “Troubleshooting FSx for ONTAP launch errors”](#)

Troubleshooting FSx for ONTAP iSCSI connectivity

When MGN migrates a server using FSx for ONTAP, the target instance must establish iSCSI sessions to the ONTAP SVM. A postboot script (`mgn_iscsi_postboot`) runs automatically to configure iSCSI connectivity with multipath redundancy when available. If no iSCSI sessions are established within the validation window (15 minutes for Linux, 25 minutes for Windows), the job fails.

Postboot log location:

- Linux: `/var/log/mgn_iscsi_postboot.log`
- Windows: `C:\Windows\Temp\mgn_iscsi_postboot.log`

[Connect to the target instance](#) via [SSM Session Manager](#) or SSH to read the log. The target instance ID is shown in the MGN console Launch status section. Look for ERROR or WARNING entries.

Common failures:

Symptom	Cause	Resolution
Package installation errors or No supported package manager found	Target instance cannot reach OS package repositories	Ensure outbound internet access (NAT gateway or internet gateway) for package downloads. See Step 6 .
Package installation fails due to missing or inaccessible repositories (SLES, RHEL, CentOS)	Instance-bound or subscription-based repository credentials are not valid on the migrated target instance	Pre-install the required iSCSI and multipath packages on the source server before migration. See the packages table in Step 6 .
Connection refused or timeout on port 3260	Security groups or NACLs not configured for iSCSI traffic	Allow TCP 3260 between the target instance and FSx for ONTAP security groups. See Step 1 .
No route to host or network unreachable	No network path between target instance and FSx for ONTAP subnets	Verify routing between the target subnet and FSx for ONTAP subnets (VPC route tables , VPC peering, or transit gateway).
Multipath-IO could not be installed	MPIO packages unavailable or Windows Desktop SKU without the feature	The migration will succeed with single-path connectivity. For full HA redundancy, ensure the instance can install MPIO (Linux: multipath-tools or device-mapper-multipath packages; Windows: Server SKU with MultiPath-IO feature available).
Phase 1 completed but Phase 2 never starts (Windows only)	Instance failed to reboot after MPIO installation or	Verify the instance is running in the EC2 console. Check Windows Event

Symptom	Cause	Resolution
	postboot service did not re-trigger	Viewer for boot errors. Retry the launch.
Log file does not exist	Postboot script did not run (conversion failure before postboot stage)	Check the MGN console for earlier errors in the launch status.

Note

If the validation reports "iSCSI connectivity established with 1 of 2 expected sessions - operating without multipath", the migration succeeds but without full HA redundancy. Verify that both FSx for ONTAP subnet endpoints are routable from the target instance.

For manual iSCSI verification, see [Mounting iSCSI LUNs on FSx for ONTAP](#).

After fixing the issue, launch a new test or cutover from the MGN console. The postboot script will run again automatically.

FSx for ONTAP replication errors

When MGN encounters storage issues while replicating data to FSx for ONTAP, replication stalls and an error is displayed in the MGN console. This section covers the common replication stall errors and how to resolve them.

Failed to start data transfer

If replication stalls with a "Failed to Start Data Transfer" error, MGN could not begin writing data to the FSx for ONTAP file system. This may be caused by a storage capacity issue.

Possible causes and resolutions:

Cause	How to verify	Resolution
File system is out of storage capacity	In the FSx console , check the file system's Storage	Increase the file system's storage capacity. For more information,

Cause	How to verify	Resolution
	capacity and Used storage metrics.	see Managing storage capacity and provisioned IOPS .
Replication volume is out of storage capacity	In the FSx console , check for volume-level errors on the replication volume.	Increase the volume size to accommodate the replicated data.

After resolving the issue, replication recovers automatically. It may take up to a few hours for the stall indicator to clear in the MGN console.

Not converging

If replication enters a "Not Converging" state, MGN is unable to keep up with changes on the source server. The rate of incoming data exceeds the rate at which MGN can write to the FSx for ONTAP file system. In addition to the general causes listed in [Common replication errors](#), the following FSx for ONTAP-specific causes may apply:

Possible causes and resolutions:

Cause	How to verify	Resolution
File system is out of storage capacity	In the FSx console , check the file system's Storage capacity and Used storage metrics.	Increase the file system's storage capacity. For more information, see Managing storage capacity and provisioned IOPS .
Throughput capacity is insufficient for the workload	In the FSx console, check the Throughput CloudWatch metrics for the file system. Look for sustained throughput near the provisioned limit.	Increase the file system's throughput capacity. You can modify throughput at any time. For more information, see Managing throughput capacity .

After resolving the issue, replication recovers automatically. It may take up to a few hours for the stall indicator to clear in the MGN console.

Storage operation timed out

If a migration operation fails with a storage operation timeout, this indicates that MGN could not complete a storage request to the FSx for ONTAP file system within the expected time. This can be caused by insufficient capacity or degraded performance.

Possible causes and resolutions:

Cause	How to verify	Resolution
File system is out of storage capacity	In the FSx console , check the file system's Storage capacity and Used storage metrics.	Increase the file system's storage capacity. For more information, see Managing storage capacity and provisioned IOPS .
Throughput capacity is insufficient for the workload	In the FSx console, check the Throughput CloudWatch metrics for the file system. Look for sustained throughput near the provisioned limit.	Increase the file system's throughput capacity. You can modify throughput at any time. For more information, see Managing throughput capacity .

After resolving the issue, replication recovers automatically. It may take up to a few hours for the stall indicator to clear in the MGN console.

Replication volume not deleted after Finalize cutover/Disconnect from service (FlexClone split blocked by backup)

Symptoms:

- After finalize cutover, the replication volume remains in the FSx for ONTAP file system and is not cleaned up by MGN.
- The MGN console may show the source server in "Cutover" state but the old replication volume persists.

Cause:

After finalize cutover, MGN creates a FlexClone from the replication volume and initiates a split to make it independent. If FSx for ONTAP automatic backups (or a manual backup) were taken on the target volume (FlexClone) after finalize cutover, the backup creates a locked snapshot that blocks the FlexClone split operation. MGN cannot complete the split until the locked snapshot is removed.

Resolution:

1. **Check if automatic backups are enabled** – In the FSx for ONTAP console, navigate to your file system and check the backup settings. If automatic backups are enabled, disable them to prevent new backups from being created on the volumes.
2. **Delete backups from target volumes** – In the FSx for ONTAP console, navigate to **Backups** and delete all backups associated with volumes matching the target_*source_server_id_timestamp* pattern. There may be more than one target volume per source server. Select each backup and choose **Actions** → **Delete backup**. This releases the locked snapshots that block the split.

 Note

Deleting the backup does not affect the target volume data. If you need to retain backup data, you can restore it to a separate volume before deleting.

3. **Wait 24 hours for cleanup** – MGN will automatically complete the FlexClone split and delete the replication volume within 24 hours. No further manual action is needed.
4. **Re-enable automatic backups** – After the replication volume has been cleaned up, re-enable automatic backups on the FSx for ONTAP file system to resume regular backup protection.

For more information, see [Managing FSx for ONTAP volumes](#) and [FSx for ONTAP backups](#).

Orphaned FSx for ONTAP target volumes (FlexClone) after launch cleanup

Symptom:

After a "Revert to Ready for testing", "Revert to Ready for cutover", or "Terminate launched instances" action, one or more FSx for ONTAP volumes prefixed with `atx_cleanup_required_` remain on the file system.

Cause:

MGN cannot delete FlexClone volumes that have active FSx for ONTAP backup relationships (SnapMirror) on them. When this occurs, MGN renames the volume with the `atx_cleanup_required_` prefix to indicate that it requires manual cleanup.

How to confirm:

1. Navigate to the MGN console → **Launch history**.
2. Find the job corresponding to the termination action.
3. Check the job event logs for the following message:

```
"N FSx for ONTAP volume(s) (prefixed 'atx_cleanup_required_') could not be deleted due to active backup relationships. Delete these volumes from the FSx console manually."
```

Finding volumes that require cleanup:

1. Open the FSx for ONTAP console → **Volumes**.
2. Filter or search for volumes with names starting with `atx_cleanup_required_`.

Resolution:

Delete the orphaned volume manually via the FSx for ONTAP console:

1. Open the FSx for ONTAP console → **Volumes**.
2. Locate the volume matching the `atx_cleanup_required_` prefix.
3. Delete the volume.

Verifying cleanup is complete:

Confirm that no volumes with the `atx_cleanup_required_` prefix remain in your FSx for ONTAP file system.

Troubleshooting FSx for ONTAP launch errors

Use the information in this section to troubleshoot launch errors specific to FSx for ONTAP migrations.

Insufficient file system capacity

Before launching a test or cutover, MGN validates that the FSx for ONTAP file system has sufficient capacity. The launch fails if predicted usage would exceed 90% of the aggregate capacity, with an error similar to:

```
Launch check failed: insufficient capacity on file system fs-0123456789abcdef0
(85% used). This job needs ~50 GB of free space. Expand storage and retry.
```

Cause

The FSx for ONTAP file system does not have enough free SSD storage capacity to create FlexClone volumes for all source servers in the launch job.

Resolution

Increase the SSD storage capacity of your FSx for ONTAP file system.

1. Open the FSx for ONTAP console at <https://console.aws.amazon.com/fsx/>, and choose **File systems**.
2. Select the file system shown in the error message.
3. On the **Summary** panel, choose **Update** next to **SSD storage capacity**.
4. Enter the new desired capacity. The minimum increase is 10% of the current capacity or 1 TiB, whichever is greater.
5. Choose **Update**.
6. Wait for the storage update to complete, then retry the launch from the MGN console.

Note

Storage capacity increases are non-disruptive. The file system remains available during the scaling operation.

Tip

During migration, the file system holds both the replica volumes (used for ongoing replication) and the cloned volumes (created at launch). Both coexist until you finalize the

cutover and delete the replica volumes. Plan your SSD capacity to accommodate both sets simultaneously, or reduce the number of source servers in a single launch job.

For more information, see [Managing SSD storage capacity and provisioned IOPS](#).

FAQ

Use this section to find answers to many frequently asked questions.

Topics

- [General questions](#)
- [Agent related FAQs](#)
- [Agentless replication related FAQs](#)
- [Replication related FAQs](#)
- [AWS related FAQs](#)
- [Does MGN work with...?](#)
- [Post-launch actions related FAQs](#)
- [Storage related FAQs](#)

General questions

This section contains answers to general questions about AWS Transform MGN.

Topics

- [Why was AWS Application Migration Service renamed to AWS Transform MGN?](#)
- [Can MGN protect or migrate physical servers?](#)
- [What data is stored on and transmitted through MGN service?](#)
- [What should I consider when replicating Active Directory?](#)
- [Does AWS Transform MGN work with LVM and RAID configurations?](#)
- [What is there to note regarding SAN/NAS support?](#)
- [Does AWS Transform MGN support Windows License migration?](#)
- [Can you perform an OS \(Operating System\) upgrade with AWS Transform MGN?](#)
- [What are the AWS Transform MGN quota limits?](#)
- [What are the Private APIs used by MGN to define actions in the IAM Policy?](#)
- [Which post-launch scripts does MGN support?](#)
- [What happens if I use a custom DNS?](#)

- [Can I use AWS Transform MGN to migrate servers from VMware Cloud on AWS \(VMC\) to Amazon EC2?](#)
- [When should I use AWS Elastic Disaster Recovery \(AWS DRS\) for migration?](#)

Why was AWS Application Migration Service renamed to AWS Transform MGN?

AWS Application Migration Service has been renamed to AWS Transform MGN. The new name reflects the close link between MGN and AWS Transform. AWS Transform uses MGN replication technology to rehost servers.

You can access AWS Transform MGN in two ways:

- Use the MGN console directly in the AWS console for a hands-on experience.
- Use the AWS Transform workflow, which automates discovery, wave planning, network setup, landing zone creation, rehosting and containerization.

At the rehosting stage, you can continue with the agentic workflow or switch to the MGN console if you prefer a more hands-on experience.

AWS Transform is also available through Kiro, Claude, Cursor, and Codex via the AWS Transform MCP server.

Can MGN protect or migrate physical servers?

Yes. MGN can migrate both virtual and physical servers. The replication process works at the OS level, so it handles both server types the same way.

What data is stored on and transmitted through MGN service?

MGN stores only configuration and log data. This data is kept in an encrypted database. Your replicated data stays in your own VPC. All data in transit is encrypted.

What should I consider when replicating Active Directory?

There are two main approaches when it comes to migrating Active Directory or domain controllers from a disaster:

1. Replicating the entire environment, including the AD server(s) – in this approach it is recommended to launch the test or cutover AD servers first, wait until it's up and running, and then launch the other test or cutover instances, to make sure the AD servers are ready to authenticate them.
2. Leaving the AD server(s) in the source environment – in this approach, the test or cutover instances will communicate back to the AD server in the source environment and will take the source server's place in the AD automatically.

In this case, it is important to conduct any tests using an isolated subnet in the AWS cloud, so to avoid having the test or cutover instances communicate into the source AD server outside of a cutover.

Does AWS Transform MGN work with LVM and RAID configurations?

Yes, AWS Transform MGN works with any such configuration.

What is there to note regarding SAN/NAS support?

If the disks are represented as block devices on the machine, as most SAN are, AWS Transform MGN will replicate them transparently, just like actual local disks.

If the disks are mounted over the network, such as an NFS share, as most NAS implementations are, the AWS Replication Agent would need to be installed on the actual NFS server to replicate the disk.

Does AWS Transform MGN support Windows License migration?

AWS Transform MGN conforms to the [Microsoft Licensing on AWS](#) guidelines.

Can you perform an OS (Operating System) upgrade with AWS Transform MGN?

Yes. AWS Transform MGN allows you to [perform an OS upgrade](#) using a predefined action. The action will clone your machine and upgrade the clone. After the upgrade, verify that the cloned machine is working well, and then you can begin using it.

What are the AWS Transform MGN quota limits?

The following are the AWS Transform MGN service quota limits:

Name	Default	Description
Concurrent jobs in progress	Each supported AWS Region: 20	Launching a test or cutover instance, or a cleanup action is considered a "job". Multiple servers launched together count as a single job. This parameter is the maximum number of Jobs that can be run concurrently. Jobs that are Completed are not counted against this quota.
Max active source servers	Each supported AWS Region: 150	The maximum number of servers that can be actively replicating at any time. For larger migrations contact Support.
Max non-archived source servers	Each supported AWS Region: 4,000	This parameter is used for agentless migrations. This is the max number of servers that can be managed by MGN, in non-archived state. This includes the servers that are actively replicating, as well as any servers whose replication has not yet started. The number of actively replicating servers is controlled by the parameter Max active source servers .
Max source servers in a single job	Each supported AWS Region: 200	Launching a test or cutover instance, or a cleanup action is considered a "Job". If you select multiple servers, and

Name	Default	Description
		perform one of these actions, they are grouped into a single job. This is the maximum number of servers that can be grouped into a single Job.
Max source servers in all jobs	Each supported AWS Region: 200	Launching a test or cutover instance, or a cleanup action is considered a "Job". This is the maximum total number of servers that can be configured in all active Jobs. Jobs that are Completed are not counted against this quota.
Max total source servers per AWS account	Each supported AWS Region: 50,000	This parameter is the maximum total servers, both active and archived, that can be migrated in a single account in each AWS Region. Servers that are deleted, are not counted against this quota.
Max concurrent jobs per source server	Each supported AWS Region: 1	Launching a test or cutover instance, or a cleanup action is considered a "Job". This is the maximum number of active Jobs, that can be configured per server. Jobs that are Completed are not counted against this quota.

You can learn about the AWS Transform MGN limits in the [AWS General Reference](#).

What are the Private APIs used by MGN to define actions in the IAM Policy?

MGN uses the following Private API resources as actions in the IAM Policy. [Learn more about Actions, resources, and condition keys for MGN.](#)

- `BatchCreateVolumeSnapshotGroupForMgn` – Grants permission to create volume snapshot group.
- `BatchDeleteSnapshotRequestForMgn` – Grants permission to batch delete snapshot request.
- `DescribeReplicationServerAssociationsForMgn` – Grants permission to describe replication server associations.
- `DescribeSnapshotRequestsForMgn` – Grants permission to describe snapshots requests.
- `GetAgentCommandForMgn` – Grants permission to get agent command.
- `GetAgentConfirmedResumeInfoForMgn` – Grants permission to get agent confirmed resume info.
- `GetAgentInstallationAssetsForMgn` – Grants permission to get agent installation assets.
- `GetAgentReplicationInfoForMgn` – Grants permission to get agent replication info.
- `GetAgentRuntimeConfigurationForMgn` – Grants permission to get agent runtime configuration.
- `GetAgentSnapshotCreditsForMgn` – Grants permission to get agent snapshots credits.
- `GetChannelCommandsForMgn` – Grants permission to get channel commands.
- `NotifyAgentAuthenticationForMgn` – Grants permission to notify agent authentication.
- `NotifyAgentConnectedForMgn` – Grants permission to notify agent is connected.
- `NotifyAgentDisconnectedForMgn` – Grants permission to notify agent is disconnected.
- `NotifyAgentReplicationProgressForMgn` – Grants permission to notify agent replication progress.
- `RegisterAgentForMgn` – Grants permission to register agent.
- `SendAgentLogsForMgn` – Grants permission to send agent logs.
- `SendAgentMetricsForMgn` – Grants permission to send agent metrics.
- `SendChannelCommandResultForMgn` – Grants permission to send channel command result.
- `SendClientLogsForMgn` – Grants permission to send client logs.
- `SendClientMetricsForMgn` – Grants permission to send client metrics.
- `UpdateAgentBacklogForMgn` – Grants permission to update agent backlog.
- `UpdateAgentConversionInfoForMgn` – Grants permission to update agent conversion info.
- `UpdateAgentReplicationInfoForMgn` – Grants permission to update agent replication info.

- `UpdateAgentReplicationProcessStateForMgn` – Grants permission to update agent replication process state.
- `UpdateAgentSourcePropertiesForMgn` – Grants permission to update agent source properties.
- `CreateVcenterClientForMgn` – Grants permission to create a vCenter client.
- `GetVcenterClientCommandsForMgn` – Grants permission to get vCenter client commands.
- `SendVcenterClientCommandResultForMgn` – Grants permission to send vCenter client command result.
- `SendVcenterClientLogsForMgn` – Grants permission to send vCenter client logs.
- `SendVcenterClientMetricsForMgn` – Grants permission to send vCenter client metrics.
- `NotifyVcenterClientStartedForMgn` – Grants permission to notify vCenter client started.
- `IssueAgentCertificateForMgn` – Grants permission to send certificate signing request.

Which post-launch scripts does MGN support?

MGN can run scripts on a launched test or cutover instance. This is done by creating the following folder on the source server and placing the scripts within that folder.

Linux: `/boot/post_launch` (any files that are marked as executable)

Windows: `C:\Program Files (x86)\AWS Replication Agent\post_launch\` (any .exe, .cmd, or .bat files)

Once you put these scripts in the above folders on the source server, the folder will be replicated to the test or cutover instance and be executed once after the instance boots for the first time.

Note

Post-launch scripts on Windows run under the Local System context. Post-launch scripts on Linux run under the 'root' user.

Uninstalling VMTools from Windows

The following script can be utilized to uninstall VMTools post migration from Windows. This is a powershell script. It needs to be wrapped by a .CMD file, as powershell scripts are not run automatically by the `post_launch`.


```
$regpath = "HKLM:\Software\Microsoft\Windows\CurrentVersion\uninstall"

Get-childItem $regpath | % {

$keypath = $_.pschildname

$key = Get-Itemproperty $regpath\$keypath

if ($key.DisplayName -match "VMware Tools") {

$VMwareToolsGUID = $keypath

}

MsiExec.exe /x $VMwareToolsGUID /qn /norestart

}
```

What happens if I use a custom DNS?

Custom DNS settings can cause issues in the replication servers.

Therefore, if you are using a custom DNS, you will need to add a TCP port 53 to the security group outbound rules, for replication and conversion servers.

Can I use AWS Transform MGN to migrate servers from VMware Cloud on AWS (VMC) to Amazon EC2?

Yes, you can. For migrations of source servers from [VMC](#) to EC2 you have two options. You can install the agentless appliance in your VMC environment, and migrate your servers using [agentless replication](#), or install the [AWS replication agent](#) on each of your source servers, and use agent-based replication for your migration.

When should I use AWS Elastic Disaster Recovery (AWS DRS) for migration?

In cases that DRS supports a feature that does not exist in MGN, DRS can be used for migration. You can install the DRS replication agent on your source servers. Following replication, you can launch recovery instances in your target environment, to complete the migration.

DRS can be used for migration, as the DRS and MGN services use shared technology for performing block level replication. Both MGN and DRS have a replication agent, for replicating servers into a staging area in AWS. MGN supports launching test and cutover instances from the staging area. DRS supports launching recovery instances from the staging area. The technology used by both of these services for launching instances in AWS is very similar. DRS also has the capability to failback to the source environment, after the source environment has recovered. This capability does not exist in MGN.

Note that you cannot install the DRS and MGN agents on the same server at the same time. If you already installed the MGN agent on a server, and want to use DRS for migration, you must uninstall the MGN agent before installing the DRS agent.

Note that there are costs associated with using the DRS service. For DRS pricing information see [AWS Elastic Disaster Recovery pricing](#).

Agent related FAQs

This section contains answers to questions about the AWS Replication Agent.

Topics

- [What does the AWS Replication Agent do?](#)
- [What kind of data is transferred between the agent and the AWS Transform MGN?](#)
- [Can a proxy server be used between the source server and the AWS Transform MGN console?](#)
- [What are the prerequisites needed to install the AWS Replication Agent?](#)
- [What ports does the AWS Replication Agent use?](#)
- [What privileges does the AWS Replication Agent require?](#)
- [Is it possible to install the agent on servers running operating systems that are not listed as supported?](#)
- [What kind of resources does the AWS Replication Agent use?](#)

- [Can AWS Transform MGN migrate containers?](#)
- [Does the AWS Replication Agent cache any data to disk?](#)
- [How is communication between the AWS Replication Agent and the AWS Transform MGN secured?](#)
- [Is it possible to change the port the AWS Replication Agent uses from TCP Port 1500 to a different port?](#)
- [How do I manually uninstall the AWS Transform MGN agent from a server?](#)
- [When do I need to reinstall the agent?](#)
- [How much bandwidth does the AWS Replication Agent consume?](#)
- [How many disks can the AWS Replication Agent replicate?](#)
- [Is it possible to add a disk to replication without a complete resync of any disks that have already been replicated?](#)
- [Is the AWS Replication Agent installed on launched test and cutover instances?](#)
- [How do temporary credentials work?](#)
- [Which Windows and Linux OSs support no-rescan upon reboot?](#)

What does the AWS Replication Agent do?

The AWS Replication Agent performs an initial block-level read of the content of any volume attached to the server and replicates it to the replication server. The agent then acts as an OS-level read filter to capture writes and synchronizes any block level modifications to the AWS Transform MGN replication server, ensuring near-zero RPO.

What kind of data is transferred between the agent and the AWS Transform MGN?

The AWS Replication Agent sends the following types of information to the Service Manager of AWS Transform MGN:

- Monitoring metrics of the agent itself
- Replication status (started, stalled, resumed)
- Backlog information
- OS and hardware information

When an Agent is installed on a source server, it collects the following information on the machine:

- Host name and ID
- List of CPUs including models and number of cores
- Amount of RAM
- Hardware and OS information
- Number of disks and their size – in Windows, disk letters; in Linux, block device names
- Machine's Private IP address

Can a proxy server be used between the source server and the AWS Transform MGN console?

Yes. The proxy is configured using an environment variable before the install.

`https_proxy=https://PROXY:PORT/`

For example: `https_proxy=https://10.0.0.1:8088/`

Make sure the proxy has a trailing forward slash.

Ensure that you have allowlisted the [MGN IPs and URLs](#) for both SSL Interception and Authentication.

Note

A web proxy cannot be used for communication between the source server and the staging area subnet where replication server launched for replication over TCP Port 1500. To use private routing for data replication, see [Data routing and throttling](#).

What are the prerequisites needed to install the AWS Replication Agent?

The installation requirements for source server depend on the type of OS that the server runs – either Linux or Windows.

Prerequisites [can be found here](#).

What ports does the AWS Replication Agent use?

The Agent uses TCP Port 443 to communicate with the Service Manager of Application Migration Service and TCP Port 1500 for replication to AWS.

What privileges does the AWS Replication Agent require?

The AWS Replication Agent installer requires root privileges or the use of the sudo command during installation. It creates an "aws-replication" group and user, and attempts to add the "aws-replication" user to the "sudoers" file to grant necessary permissions. Ensure that the user running the installation has sufficient privileges to modify the "sudoers" file. If the installation fails due to insufficient permissions, you may need to manually add the "aws-replication" user to the "sudoers" file before attempting the installation again.

Is it possible to install the agent on servers running operating systems that are not listed as supported?

The agent is designed and tested to work on the officially supported operating systems listed in the documentation. Installing the agent on other unsupported operating systems might be possible but is not recommended. Any installation or replication issues encountered when using unsupported operating systems will need to be handled through your own troubleshooting or support channels, as the AWS engineering team will be limited in their ability to assist. We advise using the agent only on supported OS versions to ensure the best experience. Refer to [Supported operating systems](#).

What kind of resources does the AWS Replication Agent use?

The AWS Replication Agent is lightweight and nondisruptive. The agent uses approximately 5% CPU and 250 MB of RAM.

Can AWS Transform MGN migrate containers?

AWS Transform MGN only supports the replication of full servers. Nevertheless, MGN replicates on a server level and therefore any containers within the selected servers will be replicated.

Does the AWS Replication Agent cache any data to disk?

AWS Transform MGN does not write any cache or do any sort of journalling to disk. The Agent holds a buffer which is large enough to map all volume's blocks ~250 MB in memory.

The agent then acts as a sort of write filter and will replicate changed blocks directly from memory to the Replication Server. In cases where the data is no longer in memory, the agent will read the block from the volume directly. This is the case where you may see backlog in the AWS Transform MGN console. The cause of this is the volume of change is greater than the bandwidth available.

How is communication between the AWS Replication Agent and the AWS Transform MGN secured?

All communication is encrypted using SSL. In addition, each Agent is assigned a key during installation which is used to encrypt all traffic. All keys are unique and are not shared across multiple agents.

Is it possible to change the port the AWS Replication Agent uses from TCP Port 1500 to a different port?

No. The AWS Transform MGN Agent can only use TCP Port 1500 for replication.

How do I manually uninstall the AWS Transform MGN agent from a server?

Follow the steps in the [the section called "Uninstalling the Agent"](#) section.

When do I need to reinstall the agent?

Typically, you need to reinstall the Agent after any major upgrade to the source server.

Linux

- Any kernel upgrade
- After adding new volumes

Windows

- Any OS upgrade (for example, Windows Server 2012 to Windows Server 2016)

Note

If you [upgrade using a post-launch action](#), an agent upgrade is not required.

- After adding new volumes

How much bandwidth does the AWS Replication Agent consume?

The AWS Replication Agent opens up to five connections and will attempt to maximize available bandwidth.

Throttling can be activated via the AWS Transform MGN console by either selecting a specific server and choosing the **Replication settings** tab or by changing the **Replication template** (in this case the change will only affect newly added servers).

How many disks can the AWS Replication Agent replicate?

The agent can replicate up to 63 disks from a single server. Ensure that the replication server instance type supports at least the number of disks being replicated.

Is it possible to add a disk to replication without a complete resync of any disks that have already been replicated?

When you are adding a disk to a source server, AWS Transform MGN will not automatically identify this disk and add it to the **Disk settings** section in the console.

The only way to get this disk to replicate is to reinstall the agent. Before reinstalling, you can note the current **Total replicated storage**. When you reinstall the agent, you will notice the value of replicated storage changes.

You will also notice an additional progress bar appear, which indicates that we are rescanning the original volumes. This is not a resync, but a scan, to verify that all the blocks on the source still match the blocks on the replication side. This process is significantly quicker than a resync, as there is no actual block data transferred, unless there is a difference. This is needed, as a reinstall results in the driver which performs the IO tracking being unloaded and reset, so we have no way of being certain of the sync status. While the rescan on the original volumes is happening, the agent is also ensuring that the initial sync of the new volume is being completed in parallel.

Is the AWS Replication Agent installed on launched test and cutover instances?

During the launch process, either upon test or cutover instance launch, the AWS Replication agent is removed from the test or cutover instance, and will not run on it.

How do temporary credentials work?

The temporary credential mechanism was developed specifically to provide an easy and secure way to install MGN Agents. The main flow of the temporary credentials' creation process relies on generating an x509 certificate per agent and then using this x509 certificate to receive temporary IAM credentials. This process uses a similar mechanism to the one used by [IAM Roles Anywhere](#).

Which Windows and Linux OSs support no-rescan upon reboot?

A shutdown (from the OS menu or CLI) of any supported Linux or Windows source server no longer causes a rescan in MGN once the source server is restarted.

A rescan means that the agent on the source server rereads all blocks on all replicated disks and transmits blocks that are different from the previously replicated data. A rescan is similar to the initial sync but is faster because only blocks that are different need to be transmitted.

Rescans can still happen following a hard reboot, crashes, or when you add or remove disks to or from the source server.

Supported OSs include:

Windows Server

- 2012r1
- 2012r2
- 2016
- 2019
- 2022
- 2025
- Windows 11

Linux

- CentOS 6–8, Stream 9, Stream 10
- Oracle 6–8
- RHEL 6–9.8, 10, 10.1, 10.2

- Rocky Linux 8–9.8, 10, 10.1, 10.2
- SLES 12 and 15
- Debian 9–11
- Ubuntu 16, 18, 20, and 22
- Amazon Linux 2 and 2023
- AlmaLinux 9.6, 9.7, 9.8, 10, 10.1, 10.2

Note

For Linux, no-rescan on reboot is supported only on environments that use initramfs.

Agentless replication related FAQs

This section contains answers to questions about agentless replication.

Topics

- [In which situations would you recommend using agentless replication \(snapshot shipping\)?](#)
- [In which situations would you recommend using agent-based replication?](#)
- [How does agentless replication work?](#)
- [Does agentless replication require installing any component in the customer's source data center?](#)
- [Is the agentless feature available in all Regions that MGN service supports?](#)
- [Does agentless replication support the same source operating systems that are supported by agent-based replication?](#)
- [Is the agentless feature supported in CloudEndure migration?](#)
- [Which virtualization environments are supported by the agentless feature?](#)
- [On which operating systems can the MGN vCenter Client be installed?](#)
- [Do I need to generate special credentials to install the MGN vCenter Client?](#)
- [What are the agentless replication prerequisites?](#)
- [How do I install the MGN vCenter Client?](#)

- [Can a proxy server be used between the source server and the AWS Transform MGN console?](#)

In which situations would you recommend using agentless replication (snapshot shipping)?

Agentless replication best serves customers whose company's security policies do not allow installing an agent on each of their source servers, or for operating systems that are only supported by agentless replication. This solution is only available for data centers using vCenter version 6.7, 7.0 and 8.0.

In which situations would you recommend using agent-based replication?

Agent-based replication is our default recommendation for all use cases, except when your company's security policies prevent you from using this method or if the OS is not supported. Using agent-based replication provides Continuous Data Replication, and ensures a cutover window of minutes. When using agentless replication, the data is transferred using snapshot shipping. Upon cutover, you may need to wait to have a fully updated snapshot, and your cutover window may be longer.

How does agentless replication work?

You can learn more about how agentless replication works and see a high-level diagram of the agentless replication framework in the [agentless replication documentation](#).

Does agentless replication require installing any component in the customer's source data center?

Yes. To use agentless replication, customers must install the MGN vCenter Client in their source data center. The client discovers the source servers and replicates their data to AWS.

Is the agentless feature available in all Regions that MGN service supports?

Yes. Both agent-based and agentless replication is supported in AWS Transform MGN in all Regions.

Does agentless replication support the same source operating systems that are supported by agent-based replication?

Agentless replication supports all of the [supported Windows operating systems](#) and [supported Linux operating systems](#) of agent-based replication.

Is the agentless feature supported in CloudEndure migration?

No. This feature is only available on AWS Transform MGN.

Which virtualization environments are supported by the agentless feature?

The agentless replication feature is available for vCenter versions 6.7, 7.0 and 8.0. Other virtualization environments are not supported.

On which operating systems can the MGN vCenter Client be installed?

The MGN vCenter Client can be installed on the following 64 bit Linux versions:

- Ubuntu 18.x+ (64 bit) - 22.04
- Amazon Linux 2
- RHEL 8.x

Do I need to generate special credentials to install the MGN vCenter Client?

Yes. To use the MGN vCenter Client, you must first generate the correct IAM credentials. Learn more in the [agentless replication documentation](#).

What are the agentless replication prerequisites?

The only prerequisite for agentless replication is to ensure that you have initialized AWS Transform MGN.

How do I install the MGN vCenter Client?

You can learn more about installing the MGN vCenter Client as well as installation requirements in the [agentless replication documentation](#).

Can a proxy server be used between the source server and the AWS Transform MGN console?

Yes. You can configure transparent proxy either by using an environment variable before the installation, or by using the `--proxy-address` flag in the Linux installer:

- Using the installer: `./aws-vcenter-client-installer-init.py --proxy-address http://PROXY:PORT/`
- Using environment variable: `export https_proxy=http://PROXY:PORT/; ./aws-vcenter-client-installer-init.py`

Make sure the proxy has a trailing forward slash (/).

Replication related FAQs

This section contains answers to questions about data replication.

Topics

- [What is the lifecycle of the snapshots and volumes automatically created during migration?](#)
- [What do Lag and Backlog mean during replication?](#)
- [What is Continuous, Block level data replication?](#)
- [What are the Replication initiation Steps?](#)
- [Is the replicated data encrypted?](#)
- [How is the Replication Server provisioned and managed in the Staging Area?](#)
- [What type of replication server is utilized in the AWS Transform MGN staging area?](#)
- [Can we set specific IP addresses for the replication server or conversion server in the AWS Transform MGN staging area?](#)
- [Does AWS Transform MGN compress data during replication?](#)
- [Are events that are generated by the AWS Transform MGN servers logged in Cloudtrail in AWS?](#)
- [How many snapshots does AWS Transform MGN create?](#)
- [Does AWS Transform MGN delete snapshots?](#)

- [How much capacity is allocated to the staging area?](#)
- [Why is 0.0.0.0:1500 added to inbound rules in the staging area?](#)
- [Can AWS Transform MGN replicate Oracle ASM?](#)
- [How long does a rescan take?](#)
- [How can I control the bandwidth used for replication?](#)
- [Are migrations performed by Application Migration Service crash consistent?](#)
- [How can I perform an SSL connectivity and bandwidth test?](#)
- [Why are MGN replication and conversion servers changing to Amazon Linux 2023?](#)

What is the lifecycle of the snapshots and volumes automatically created during migration?

For each source block device, MGN creates a corresponding EBS volume. If the agent on the source machine cannot send data to a volume, MGN creates a replacement volume. The old volume might remain for approximately 10 minutes after the replacement volume comes online. Volume replacement occurs only when there is an agent communication failure, which typically results from an unstable network connection.

MGN takes regular EBS snapshots to use incremental snapshot capabilities. Frequent snapshots reduce the time required to create each snapshot, which means that test or cutover instance launches are not delayed while waiting for EBS snapshots to complete. MGN retains 5–6 snapshots per volume to ensure that at least one completed snapshot is available at launch time. EBS snapshot creation has no SLA and can be delayed. Snapshot creation can also fail independently of the API call. MGN retains multiple snapshots to provide redundancy if the most recent snapshot fails.

What do Lag and Backlog mean during replication?

During replication you may see a server falls out of Continuous Data Protection (CDP) mode. This may occur for various reasons, typically related to the network throughput or interruption.

- **Lag** – The amount of time since the server was last in CDP mode.
- **Backlog** – The amount of data that was written to the disk and still needs to be replicated to reach CDP mode.
- **ETA** – The estimated time remaining to return to CDP.

What is Continuous, Block level data replication?

During continuous block-level replication, the replication agent continuously monitors disk input/output (IO) activity on the protected disks. It then replicates all WRITE operations, which involve modifying or adding data, to the Replication Server, ensuring that the data is duplicated and kept up-to-date on the replication server.

What are the Replication initiation Steps?

The following replication steps are involved in the automatic creation of the replication server in the staging area and initiation of data transfer over TCP port 1500:

- **Create security groups** - Creating EC2 security groups with inbound TCP port 1500 allowed. This security group will be attached to replication server.
- **Launch Replication Server** - AWS EC2 instance is launched based on the replication configuration set.
- **Boot Replication Server** - The EC2 instance completes boot process which will now function as a replication server.
- **Authenticate with service** - Using the user data scripts and the EC2 instance configuration, the instance (replication server) will authenticate with AWS Transform MGN using service/vpc endpoint.
- **Download replication software** - The Replication Server downloads replication software from S3. This replication software will write the incoming replicated data to the Replication Server disks.
- **Create staging disks** - The Replication Server creates replication disks to store the incoming replicated data. The number of replication disks that are created depends on the size of the replicated data.
- **Attach the staging disks** – In the previous step, the replication disks were created independently, without being attached to a specific Replication Server. Now, they are attached to a Replication Server.
- *Note: When using FSx for ONTAP as the target storage type, volumes are created on the FSx for ONTAP file system instead of Amazon EBS.*
- **Pair the Replication Server with AWS Replication Agent** – Until this stage, the AWS Transform MGN managed the communication between the Agent on the Source infrastructure and the Replication Server on the Target infrastructure. Now, the AWS Transform MGN knows that all the initiation steps have been completed successfully. Therefore, it provides the Agent and the

Replication Server information about each other, so that they could start communicating with each other directly.

- Connect AWS Replication Agent to the Replication Server – The Agent and the Replication Server begin communicating with each other directly.
- Start data transfer - Replication begins from source server to Replication Server over TCP 1500.

Is the replicated data encrypted?

AWS Transform MGN encrypts all the data in transit.

How is the Replication Server provisioned and managed in the Staging Area?

AWS Transform MGN provisions the Replication Server(s) and automatically manages the addition and removal of the servers as necessary. The AWS Transform MGN automatically replaces the replication server on a periodic basis. This ensures the replication server is using the latest public Amazon Machine Image (AMI) available for MGN. The latest AMI contains the most up-to-date patches, security updates, and bug fixes. Regularly updating the replication server AMIs allows MGN to maintain the replication servers with the newest components and protections. Customers do not need to take any action to receive new replication server AMIs. MGN will automatically replace the existing replication server as needed when new AMIs are released. This helps keep the replication process secure and up-to-date without any additional effort from users.

What type of replication server is utilized in the AWS Transform MGN staging area?

AWS Transform MGN provisions a t3.Small server. The typical ratio of volumes to replication servers is 15:1.

Can we set specific IP addresses for the replication server or conversion server in the AWS Transform MGN staging area?

No, you cannot specify or assign static IP addresses for the replication server or conversion server in MGN. These servers are managed and maintained by the MGN service itself. The IP addresses for these servers are dynamically assigned from the available pool of IP addresses in the Virtual Private Cloud (VPC) that you specify during the MGN setup process. It's important to note that

while you cannot control the specific IP addresses assigned to these servers, you can control the VPC and subnet in which they are deployed. This allows you to configure network access controls and security policies according to your organizational requirements.

Does AWS Transform MGN compress data during replication?

Yes, AWS Transform MGN uses LZ4 compression during transit resulting in 60–70% compression depending on the type of data.

Are events that are generated by the AWS Transform MGN servers logged in Cloudtrail in AWS?

Yes, AWS Transform MGN generates standard AWS API calls that are visible in CloudTrail.

How many snapshots does AWS Transform MGN create?

5–7 for each disk. Frequency and exact number depend on various factors, such as change rate on the source server and network stability.

There is currently no mechanism for users to adjust the frequency and number of snapshots.

Does AWS Transform MGN delete snapshots?

AWS Transform MGN automatically deletes snapshots that are no longer used (such as those left over after source servers have been removed from the AWS Transform MGN console).

How much capacity is allocated to the staging area?

A volume is created for each volume in the source infrastructure of the same size.

Why is 0.0.0.0:1500 added to inbound rules in the staging area?

AWS Transform MGN uses TCP Port 1500 for replication between the Source Agents and the replication server. The connection is open for all IPs and can be managed by ACLs or network controls to limit inbound IPs.

Can AWS Transform MGN replicate Oracle ASM?

Replication of Oracle with ASM is supported. AWS Transform MGN also fully supports the use of the Oracle ASM Filter Driver.

How long does a rescan take?

The rescan time will vary depending on the size of the source disks. The time depends on the performance of the disks (linear read), staging area disk performance, and the rate of write operations on the source server (which are sent in parallel with the rescan). The rescan is functioning normally as long as it is moving forward and is not "stuck".

How can I control the bandwidth used for replication?

We recommend you do not limit the bandwidth used for replication. If you have business reasons to require replication to use less bandwidth, you can temporarily use the throttling feature to limit the bandwidth. This will cause lag, that will automatically recover once you remove the throttling.

Are migrations performed by Application Migration Service crash consistent?

Yes, MGN's services are crash-consistent. This means that the data retrieved by MGN when the server becomes available is the data on the server at the moment before it shut down.

How can I perform an SSL connectivity and bandwidth test?

Note

This tool is designed for AWS only.

You can use our SSL bandwidth tool to check for replication bandwidth availability.

1. In your target region, launch a c5.large test server using the public AMI named CE-ssl-speedtest.
2. Select the same subnet as the subnet used in the replication settings of your source machine.
3. Make sure that the security group allows TCP Port 1500 inbound access.
4. On the source machine, browse to `https://{test_server_ip}:1500/speedtest`
5. Choose **Start**.

Note

- Browse to the web page using the test server **public** or **private** IP according to what you defined in your replication settings.
- The following are the AMI details per region:
 - ami-00b38c08ab3506ea7 – US East (N. Virginia)
 - ami-0bd8423a4d80563fc – US East (Ohio)
 - ami-00b7159e9c985a8da – US West (N. California)
 - ami-033a4924b13126a7b – US West (Oregon)
 - ami-0bf60b09675c8d9b6 – Africa (Cape Town)
 - ami-0f01375b50763621b – Asia Pacific (Hong Kong)
 - ami-0b1aeb50834102c18 – Asia Pacific (Mumbai)
 - ami-044fa8034a31d7578 – Asia Pacific (Tokyo)
 - ami-08b042df0d4c458ea – Asia Pacific (Seoul)
 - ami-0971e46306691cd68 – Asia Pacific (Osaka)
 - ami-0afd42552b236f9dd – Asia Pacific (Singapore)
 - ami-04e7cc6b5d9e8ffa1 – Asia Pacific (Sydney)
 - ami-02f31943dfd88549d – Asia Pacific (Jakarta)
 - ami-033db317ada5abd55 – Asia Pacific (Melbourne)
 - ami-01c24408802db503d – Canada (Central)
 - ami-0b8643189a66159c9 – Europe (Stockholm)
 - ami-0dd5a09d2ae8f46b3 – Europe (Ireland)
 - ami-097fb47f3a1c2bf7e – Europe (London)
 - ami-0a3f9008725d0b4d1 – Europe (Paris)
 - ami-0c65965703bb0e541 – Europe (Milan)
 - ami-01b6fcc2337f6420d – Europe (Spain)
 - ami-07b7defb87a46bb48 – Europe (Frankfurt)
 - ami-01b3e93b3ac0e1340 – Europe (Zurich)
 - ami-016edc078b48f370b – Israel (Tel Aviv)
 - ami-0c90e298af7a2e563 – Middle East (Bahrain)
 - ami-0f7c14e62ef760768 – Middle East (UAE)

- ami-0edd5ecfc56804583 – South America (São Paulo)
- ami-00eb76deb85085b3e – AWS GovCloud (US-West)
- ami-0ba277e7ced412965 – AWS GovCloud (US-East)
- Ensure that the security groups are configured to permit connectivity on inbound port 1500.

Why are MGN replication and conversion servers changing to Amazon Linux 2023?

Beginning August 15, 2026, MGN will use Amazon Linux 2023 (AL2023) for its replication and conversion servers because [Amazon Linux 2 \(AL2\) reaches end of support on June 30, 2026](#).

What happens to in-progress jobs: If you have a replication or conversion job in progress on August 15, 2026, your existing AL2 instance continues running until the end of its 30-day lifecycle, and MGN replaces it automatically. If you have no job running on August 15, 2026, the next job you start immediately uses AL2023.

New S3 bucket dependency: AL2023 introduces a new S3 bucket dependency for package management that was not required by AL2-based servers. MGN replication and conversion servers connect to the following AL2023 package repository. In the following URL, replace *region* with the AWS Region code where your staging area is configured (for example, us-east-1).

```
https://al2023-repos-region-de612dc2.s3.dualstack.region.amazonaws.com
```

Review your configuration before August 15, 2026 if any of the following apply:

1. **Route 53 DNS Firewall, network-level blocklists, or outbound DNS restrictions.** Verify that `al2023-repos-region-de612dc2.s3.dualstack.region.amazonaws.com` is not blocked. If you use a blocklist approach, check for wildcard rules (for example, `*.s3.dualstack.region.amazonaws.com`) that may catch this domain. If you use an allowlist approach, add this domain explicitly. For more information, see [Required connectivity settings](#) and [AL2023 repository configuration](#).
2. **Isolated subnets with a restrictive S3 VPC Gateway Endpoint policy.** Add `arn:aws:s3:::al2023-repos-region-de612dc2/*` to your endpoint policy to allow the replication server to download AL2023 packages.

3. **Security or compliance policies that require pre-approval of new operating systems.** If your policy requires explicit approval of new operating systems or AMI IDs before they run in your account (for example, OS allowlisting, change management controls, or regulatory requirements), plan accordingly before August 15, 2026.

AWS related FAQs

This section contains answers to questions about AWS and AWS Transform MGN.

Topics

- [What does the AWS Transform MGN Machine Conversion Server do?](#)
- [What boot modes are supported by the AWS Transform MGN?](#)
- [How can we encrypt an unencrypted AWS Transform MGN base snapshot?](#)
- [How do I change the server AMI on AWS after Migration?](#)
- [Which AWS services are automatically installed when launching a test or cutover instance?](#)
- [How long does it take to copy a disk from the AWS Transform MGN staging area to production?](#)
- [What are the differences between conversion servers and replication servers?](#)
- [Can I prevent AWS Transform MGN from cleaning up test instance resources in AWS?](#)
- [Why are my Windows server disks read-only after launching the test or cutover instance?](#)
- [What impacts the conversion and boot time of test and cutover instances?](#)
- [Why do I observe EBS volume performance issues while using test or cutover instances?](#)
- [What are the Amazon EBS volume limits for AWS Transform MGN?](#)
- [How is the AWS Licensing Model Tenancy chosen for AWS Transform MGN?](#)
- [How does AWS Transform MGN interact with Interface VPC Endpoints?](#)
- [How do I use MGN with CloudWatch and EventBridge dashboards?](#)

What does the AWS Transform MGN Machine Conversion Server do?

The machine conversion server converts the disks to boot and run on AWS.

Specifically, the machine conversion server makes bootloader changes, injects hypervisor drivers and installs cloud tools.

What boot modes are supported by the AWS Transform MGN?

The agent supports systems using either BIOS (Basic Input/Output System) or UEFI (Unified Extensible Firmware Interface) boot modes. BIOS is the traditional boot mode that initializes hardware and bootstraps the operating system. UEFI is a more modern boot firmware that provides additional boot configurations and security features over BIOS. Both boot modes are fully supported by the agent, giving users flexibility to choose the mode that best fits their systems and requirements. Users can install the agent on servers using either UEFI or legacy BIOS firmware.

How can we encrypt an unencrypted AWS Transform MGN base snapshot?

The encryption status of AWS Transform MGN base snapshots is determined by the default EBS (Elastic Block Store) encryption setting for the respective AWS region. Encryption Scenarios:

- **Default EBS Encryption Enabled:**

If the default EBS encryption is enabled for the region, the base snapshots created by MGN will be encrypted.

- **Default EBS Encryption Disabled:**

If the default EBS encryption is not enabled, the base snapshots will be unencrypted.

- **Encrypting Existing Unencrypted Base Snapshots -**

To encrypt an existing unencrypted base snapshot, follow these steps:

1. Delete the unencrypted base snapshot from the snapshots console.
2. Enable default EBS encryption for the AWS region where the MGN source environment is located.
3. Initiate a new test or cutover migration in MGN. During this process, MGN will create a new encrypted base snapshot based on the default EBS encryption setting for the region.

Note

Enabling default EBS encryption at the region level will encrypt all newly created EBS volumes and snapshots in that region.

How do I change the server AMI on AWS after Migration?

After the machine has been launched by AWS Transform MGN switching the AMI can be done by launching a vanilla machine from the required AMI, stopping that machine, detaching all the disks (including the root) and then attaching the disks from the test or cutover instance created by AWS Transform MGN.

Which AWS services are automatically installed when launching a test or cutover instance?

AWS Transform MGN automatically installs EC2Config. After installation, EC2Config automatically installs the SSM EC2 Configuration Service.

CloudWatch, AWS PowerShell or CLI are not automatically installed. This can be done by combining the AWS Transform MGN APIs and the AWS APIs – you can use the AWS Transform MGN APIs to determine the EC2 instance IDs of the machines and then use AWS API/CLI to turn on the detailed monitoring. An alternative approach would be to do it via AWS API only based on the tags you associate with the machine. A third approach would be to do so from the post-launch script.

AWS Transform MGN installs EC2Launch (Windows 2016 only). You will need to configure EC2Launch based on [these specific requirements](#). This configuration step needs to be performed post Migration using the wizard in C:\ProgramData\Amazon\EC2-Windows\Launch\Settings\Ec2LaunchSettings.exe on the test or cutover instance.

How long does it take to copy a disk from the AWS Transform MGN staging area to production?

AWS Transform MGN uses internal cloud provider snapshots. This process typically takes less than a minute and the size of the volume does not impact the time.

What are the differences between conversion servers and replication servers?

Replication servers run on Linux and conversion servers (for Windows machines) run on Windows.

The conversion is done by AWS Transform MGN automatically bringing up a vanilla Windows conversion server machines in the same subnet with the replication servers as part of the launch job.

Both conversion and replication servers have public IPs.

The conversion servers will use the same security groups as the Replication Server.

The conversion server must be able to access the MGN's service manager.

The conversion server machines, just like the Replication servers are managed automatically by AWS Transform MGN. Any attempt to disrupt their automated functionality will result in failed conversions.

Can I prevent AWS Transform MGN from cleaning up test instance resources in AWS?

AWS Transform MGN will, by default, remove any resources created during the test process either when requested by the user or when a new Test instance is launched.

To prevent this in AWS, you can [activate Termination Protection](#) for the test or cutover instance, and the resources will not be removed upon a new instance launch.

Why are my Windows server disks read-only after launching the test or cutover instance?

When launching test or cutover instances Windows Server may boot with all the disks as read-only.

This is a common issue that occurs when detaching and attaching data disks. This issue can be resolved using steps in [this Microsoft TechNet article](#).

What impacts the conversion and boot time of test and cutover instances?

Before launching the test or cutover instance, AWS Transform MGN goes through a machine conversion server process on the boot volume. The conversion process is fairly quick.

While the actual conversion process itself is quick, the time to boot the test or cutover instance varies depending on many factors unrelated to any AWS Transform MGN processes. Some of these are controllable and should be taken into account when recovery or cutover times are of importance.

- Operating system – The amount of time required to boot the operating system is dependent on the OS itself. While Linux servers typically boot quickly, Windows servers may take additional

time, due to the nature of the Windows OS. If opportunity permits, test the boot time of the source server. If Linux OS takes a long time to boot ensure to check that dhclient (Dynamic Host Configuration Protocol Client) is installed and running on the system so it can pull an IP.

- **Scheduled Windows Updates** – If the Windows server has pending patches, ensure those are installed before launching the test or cutover instance. If pending patches remain, the boot time in the cloud may be severely impacted as the patch process may start upon the initial boot.
- **Boot volume type** – Depending on services/applications, boot time may be impacted by disk performance. It is recommended that boot volumes be tested with a higher performance SSD and even by provisioning IOPs to ensure throughput. This may be more critical during the first initial boot of the server in the cloud, as all initial settings are applied. In many cases, the boot volume type may be scaled back after the initial boot and should be tested.

Note

The first boot of Windows machines on AWS may take up to 45 minutes due to Windows adjusting to the AWS virtual hardware.

Why do I observe EBS volume performance issues while using test or cutover instances?

Note

This applies only when using Amazon EBS as the target storage type.

The EBS volumes attached to the test or cutover instances are created from snapshots of converted volumes. For any volume type that was created from snapshots, the storage blocks are pulled down from Amazon S3 and written to the volume before you access them. This process might take significant time and varies based on the EBS volume type. For additional details and EBS initialization options, refer to [Initialize Amazon EBS volumes](#)

What are the Amazon EBS volume limits for AWS Transform MGN?

AWS Transform MGN does not impose its own limits on Amazon EBS volume size or the number of Amazon EBS volumes per instance. The limits for maximum volume size, number of volumes per

instance, and other Amazon EBS constraints are governed by Amazon EBS itself. For current limits, see [Amazon EBS service quotas](#) and [Instance volume limits](#).

How is the AWS Licensing Model Tenancy chosen for AWS Transform MGN?

AWS Transform MGN conforms to the [Microsoft Licensing on AWS](#) guidelines.

How does AWS Transform MGN interact with Interface VPC Endpoints?

If you use Amazon Virtual Private Cloud (Amazon VPC) to host your AWS resources, you can establish a private connection between your VPC and AWS Transform MGN. You can use this connection to allow AWS Transform MGN to communicate with your resources on your VPC without going through the public internet.

Amazon VPC is an AWS service that you can use to launch AWS resources in a virtual network that you define. With a VPC, you have control over your network settings, such as the IP address range, subnets, route tables, and network gateways. With VPC endpoints, the routing between the VPC and AWS services is handled by the AWS network, and you can use IAM policies to control access to service resources.

To connect your VPC to AWS Transform MGN, you define an *interface VPC endpoint* for AWS Transform MGN. An interface endpoint is an elastic network interface with a private IP address that serves as an entry point for traffic destined to a supported AWS service. The endpoint provides reliable, scalable connectivity to AWS Transform MGN without requiring an internet gateway, network address translation (NAT) instance, or VPN connection. For more information, see [What is Amazon VPC](#) in the *Amazon VPC User Guide*.

Interface VPC endpoints are powered by AWS PrivateLink, an AWS technology that allows private communication between AWS services using an elastic network interface with private IP addresses. For more information, see [AWS PrivateLink](#).


For more information, see [Getting Started](#) in the *Amazon VPC User Guide*.


How do I use MGN with CloudWatch and EventBridge dashboards?

You can monitor AWS Transform MGN using CloudWatch, which collects raw data and processes it into readable, near real-time metrics. AWS Transform MGN sends events to Amazon EventBridge

whenever a source server launch has completed, a source server reaches the `READY_FOR_TEST` lifecycle state for the first time, and when the data replication state becomes stalled or when the data replication state is no longer Stalled. You can use EventBridge and these events to write rules that take actions, such as notifying you, when a relevant event occurs.

You can see MGN in CloudWatch automatic dashboards:





MGN events can be selected when defining a rule from the EventBridge console:

The screenshot shows the AWS EventBridge console 'Event pattern' form. The 'Event source' dropdown is set to 'AWS services'. The search results for 'MGN' are displayed, including 'MGN Source Server Data Replication Stalled Change', 'MGN Source Server Launch Result', and 'MGN Source Server Lifecycle State Change'. The 'Event pattern' field is empty, with a placeholder text 'Enter the event JSON'. At the bottom, there are buttons for 'Copy', 'Test pattern', and 'Edit pattern'.

[Learn more about monitoring MGN.](#)

Does MGN work with...?

This section contains answers to questions about what AWS Transform MGN works with.

Does MGN work with Microsoft Windows Failover Clustering?

Yes.

Does MGN work with Bitlocker encryption?

AWS Transform MGN does not support OS-based disk encryption features such as BitLocker. These should be deactivated before using MGN.

Does MGN work with FSx for ONTAP?

Yes. MGN can replicate data volumes to an FSx for ONTAP file system. The following describes how data flows during replication and launch:

- **During replication:** The AWS Replication Agent on your source servers sends data to MGN replication servers. When FSx for ONTAP is configured as the target storage type, MGN writes the replicated data volumes to the FSx for ONTAP file system via iSCSI.
- **At launch and cutover:** The target EC2 instances connect to the FSx for ONTAP file system via iSCSI to access their data volumes. The boot volume always remains on Amazon EBS.
- **Network:** Security groups must allow iSCSI traffic (port 3260) between MGN instances and the FSx for ONTAP file system.
- **Authentication:** Communication is secured using certificate-based authentication. You store the certificates in AWS Secrets Manager, and MGN uses them to authenticate with the FSx for ONTAP file system.

Key points:

- Agent-based replication only – FSx for ONTAP is supported only with agent-based replication.
- Per-server configuration – You can use FSx for ONTAP for some servers and Amazon EBS for others in the same wave.
- Boot volume – Always stored on Amazon EBS regardless of target storage type.

- Up to 5 FSx for ONTAP file systems per account concurrently.

For full setup instructions, see [FSx for ONTAP configuration guide](#).

Post-launch actions related FAQs

This section contains answers to questions about post-launch actions.

Topics

- [What operating systems does the post-launch actions framework support?](#)
- [What version of AWS Systems Manager Agent will be installed on my instance?](#)
- [Why is the AWS Systems Manager Agent not executing my post launch actions?](#)

What operating systems does the post-launch actions framework support?

Verify that your operating systems [are supported by AWS Systems Manager](#).

What version of AWS Systems Manager Agent will be installed on my instance?

AWS Transform MGN uses the latest [AWS Systems Manager Agent](#) version available in your instance's region.

Why is the AWS Systems Manager Agent not executing my post launch actions?

- By default, [AWS Systems Manager](#) doesn't have permission to perform actions on your instances. Grant access by using an AWS Identity and Access Management (IAM) instance profile. You can create an instance profile for AWS Systems Manager by attaching one or more IAM policies that define the necessary permissions to a new role or to a role you already created. You can use the managed policy AmazonSSMManagedInstanceCore which allows an instance to use AWS Systems Manager service core functionality or create a custom policy. For more information, see [Create an IAM instance profile for AWS Systems Manager](#).
- The instances you connect to must also allow HTTPS (port 443) outbound traffic to the following endpoints:

```
ec2messages.<REGION>.amazonaws.com  
ssm.<REGION>.amazonaws.com  
ssmmessages.<REGION>.amazonaws.com
```

You can connect to the required endpoints by using interface endpoints. For more information, see [Creating VPC endpoints for AWS Systems Manager](#).

Alternatively, you can use public IP addresses for communication between your instances and the internet.

- Another reason might be that the managed instance has limited available CPU or memory resources. Although your instance might otherwise be functional, if the instance doesn't have enough available resources, you can't establish a session. For more information, see [Troubleshooting an unreachable instance](#).

Storage related FAQs

This section contains answers to frequently asked questions about storage options in AWS Transform MGN.

What target storage types does MGN support?

MGN supports two target storage types:

- **Amazon EBS** – The default storage type. Source server disks are replicated to Amazon EBS volumes and launched as standard Amazon EC2 instances with Amazon EBS attached storage.
- **Amazon FSx for NetApp ONTAP** – Data volumes are replicated to an FSx for ONTAP file system and attached to the target instance via iSCSI. The boot volume always remains on Amazon EBS.

You configure the target storage type in the replication template or per-server replication settings.

Can I use FSx for ONTAP for some servers and Amazon EBS for others?

Yes. The target storage type is configured at the replication template level and can be overridden per source server. You can migrate some servers to FSx for ONTAP and others to Amazon EBS within the same account and region.

Do I need to pre-install iSCSI packages on my source server?

MGN automatically installs iSCSI initiator and multipath packages on the target instance during the post-boot conversion. However, you must pre-install these packages on the source server if:

- The target instance does not have network access to OS package repositories (air-gapped environments or private subnets without a NAT gateway).
- The operating system uses subscription-based repositories (SUSE SLES, RHEL, CentOS) where credentials are tied to the source instance.

For the required packages by operating system, see the [Supported Linux operating systems](#) table and the [FSx ONTAP configuration guide](#).

What happens to FSx replication volumes after finalize cutover?

After finalize cutover, MGN creates a FlexClone of the replication volume, splits it to make it independent, and then deletes the original replication volume. The target volume (FlexClone) remains attached to your cutover instance. This process is automatic and typically completes within minutes, though it may take longer for large volumes.

I initialized MGN before FSx for ONTAP support was available. Do I need to do anything?

Yes. You must reinitialize MGN to create the required AWS managed roles (AWSApplicationMigrationFsxProxyRole and AWSApplicationMigrationFsxProxyLinkRole). In the MGN console, navigate to [Getting started](#) and choose **Reinitialize**.

Release notes

September 2026

- Added support for local snapshots in the following Local Zones:
 - Hanoi, Vietnam (ap-southeast-1-han-1a)
 - Athens, Greece (eu-central-1-ath-1a)
 - Istanbul, Türkiye (eu-central-1-ist-1a)

For more information, see [Store snapshots in AWS Local Zone](#).

- You can now configure the following Amazon EBS volume settings: **Volume initialization rate** and **Delete on termination** through the AWS Transform MGN account launch template. You can set the configuration in the template under the **Default EC2 launch template** section. These parameters can now be configured at the account level for new servers, or at the individual server level via their EC2 launch template. For more information, see this section in the [Launch template](#).

August 2026

- Amazon FSx for NetApp ONTAP support as a target storage type is now generally available. You can migrate block storage workloads directly to FSx for ONTAP as part of the same migration wave, without intermediate storage or separate migration tools. For setup instructions, see [FSx for ONTAP configuration](#).
- Added support for the following operating systems:
 - Amazon Linux 2023 with kernel 6.18
 - Oracle Linux 9.6
 - RHEL 10.2
 - Rocky Linux 10.2
 - AlmaLinux 10.2
- The user-provided ID for a source server is now case insensitive. For more information, see [Edit server](#).
- Resolved an issue that caused the MS-SQL license conversion post-launch action to fail for SQL Server named instances. For more information, see [Convert MS-SQL license](#).

- The inventory import and export feature now supports additional parameters, including replication settings, launch settings, and post-launch actions settings. This expands the range of server settings you can update in bulk across your inventory. For the full list of supported parameters, see [Inventory Import parameters](#).
- MGN connector setup has been simplified. The MGN console can now automatically create the required IAM roles for the connector, in an individual account or across all member accounts of your AWS Organization. The **MGNConnectorInstallerRole** is no longer required; its permissions are now included in the **AWSApplicationMigrationConnectorManagementRole**. For more information, see [Set up the MGN Connector](#).

July 2026

Added support for the following operating systems:

- RHEL 9.8
- Rocky Linux 9.8
- AlmaLinux 9.8

June 2026

- AWS Application Migration Service has been rebranded to AWS Transform MGN. All MGN capabilities remain unchanged. Same APIs, same replication engine, same compliance certifications (FedRAMP High, HIPAA, PCI DSS, ISO, SOC 1/2/3).
- MGN now supports Amazon FSx for NetApp ONTAP as a target storage type (Public Preview). You can choose between Amazon EBS and FSx for ONTAP based on your workload requirements. For more information, see [FSx for ONTAP configuration](#).
- Added support for the following operating systems:
 - Windows 11
 - Oracle Linux 9.7 and 10.1
 - CentOS Stream 10
- Added no-rescan upon reboot support for Windows Server 2025, Windows 11, CentOS Stream 9, and CentOS Stream 10. For more information, see [Supported operating systems for no-rescan upon reboot](#).

- You can now update the user-provided ID for a source server. For more information, see [Edit server](#).
- You can now delete a source server from the service. For more information, see [Actions menu](#).

May 2026

- Updated network requirements to include the Amazon Linux 2023 (AL2023) package repository bucket for staging area subnet connectivity. This applies to both firewall-restricted environments (allowlist the dual-stack Amazon S3 URL) and isolated subnets without internet access (Amazon S3 VPC gateway endpoint policy). For more information, see [Communication between the staging area subnet and S3](#).
- Added support for CentOS Stream 9 operating system.
- Added support for additional Amazon EBS storage parameters in the EC2 launch template. You can now configure these parameters for each source server through its EC2 launch template: **Delete on termination**, **Volume initialization rate**, **KMS key**, and **EBS card index**. For more information, see [Full launch template setting review](#).

April 2026

- Service launch in these regions: *Asia Pacific (New Zealand)*, *Asia Pacific (Taipei)*, *Canada West (Calgary)*, and *Mexico (Central)* Regions.

March 2026

- MGN now supports storing EBS snapshots locally within AWS Local Zones that support this feature. For more information, see [Store snapshots in AWS Local Zone](#).
- Added support for the following operating systems:
 - RHEL 9.7 and 10.1
 - Rocky Linux 9.7, 10, and 10.1
 - AlmaLinux 9.6, 9.7, 10, and 10.1
 - SUSE Linux Enterprise Server (SLES) 15 service packs 6 and 7

December 2025

- You can now use MGN over Internet Protocol version 6 (IPv6) for both service communication and application migrations.
- Added support for the following operating systems:
 - Rocky Linux 9–9.6
 - RHEL 10

October 2025

- Added support for RHEL 9.6 operating system.

September 2025

- Specify your operating system licensing approach (BYOL / LI) and tenancy when importing inventory. Learn more in [the section called “Inventory Import parameters”](#).
- Added support for Windows Server 2025 operating system.

August 2025

- The post-launch actions *Enable Refactor Spaces* and *App2Container for Replatforming* are no longer supported.

July 2025

- You can now reset launch template values during the inventory import process. Learn more in [Editing your configuration](#).
- Service launch in these regions: *Asia Pacific (Thailand)* , and *Asia Pacific (Malaysia)* Regions.

April 2025

- MGN is now authorized for Department of Defense Cloud Computing Security Requirements Guide Impact Levels 4 and 5 (DoD CC SRG IL4 and IL5) in the AWS GovCloud (US-East and US-West) Regions.

February 2025

- Enabled tagging of network interfaces during RunInstances.

Updated the AWSApplicationMigrationFullAccess policy to support tagging network interface during runInstance. If you're managing your own policy, you must include a statement allowing `ec2:CreateTags` on `arn:aws:ec2:*:*:network-interface/*` with a condition of `"ec2:CreateAction": ["RunInstances"]`.

- Added support for RHEL 9.5 operating system.

January 2025

- Enabled tagging of network interfaces during RunInstances.
- Changed the definition of the **mgn:region** import parameter, used in the [Import feature](#). Providing a region other than the console region results in an error.

October 2024

- Added support for Oracle 9.0-9.4 operating system.

September 2024

- Introduced a new predefined post-launch action: TrendMicro. [Learn more about the TrendMicro action](#).

August 2024

- Added support for updating AWS credentials for agentless replication. [Learn more about updating vCenter & AWS credentials for agentless replication.](#)

July 2024

- Introduced a new predefined post-launch action: New Relic. [Learn more about the New Relic action.](#)
- You can use MGN with workloads that require FedRAMP High categorization level in the AWS GovCloud (US-East and US-West) Regions.

June 2024

- Added support for deploying AWS Replication Agent on a secured network in the Europe (Spain), Europe (Zurich), Middle East (UAE), Asia Pacific (Hyderabad), Asia Pacific (Osaka) and Asia Pacific (Melbourne) regions. [Learn about installing the agent on a secured network.](#)
- Added support for encrypting post-launch action parameters. Learn about [post launch action.](#)
- [AWS managed policy updates](#) - Updated the AWSApplicationMigrationFullAccess policy to support SecureString parameter type in SSM Parameters Store for post-migration framework actions.
- Added support for migrating servers with Kernel versions up to 6.8.
- Added support for Ubuntu LTS 24.04 operating system.
- Introduced a new predefined post-launch action: Dynatrace. [Learn more about Dynatrace action.](#)

May 2024

- Added support for deploying AWS Replication Agent on a secured network in the Israel (Tel Aviv) region. [Learn about installing the agent on a secured network.](#)

March 2024

- Added support for migration of Linux servers retaining boot mode UEFI.

- Added support for the following operating systems:
 - Rocky Linux 9.0
 - SUSE Linux Enterprise Server 15 service packs 4 and 5
- Added support for migrating servers with Kernel versions up to 6.5.

January 2024

- Added support for agentless replication on VMware vCenter version 8. [Learn about agentless replication.](#)

December 2023

- Added support for the MGN connector to communicate with Windows servers over HTTP and to authenticate with Linux servers using a password. [Learn more about MGN connector actions.](#)
- [AWS managed policy updates](#) – Created a new revision to support MGN in AWS GovCloud and added Statement ID (SID) to a managed policy statement: AWSApplicationMigrationServiceEc2InstancePolicy.
- Added support for deploying AWS Replication Agent on a secured network in the Asia Pacific (Jakarta) Region. [Learn about installing the agent on a secured network.](#)

November 2023

- Introduced a new predefined post-launch action: App2Container for Replatforming.

September 2023

- Introduced **MGN connector**, a feature that helps automate the agent installation on source servers. [Learn more here.](#)
- Display tags as columns in the source servers, applications and waves table in the console.
- Added support for Amazon Linux 2023 operating system.
- Added support for kernel versions up to 6.1.
- Added support for using agentless replication with a proxy server. For more information see [agentless replication installation instructions.](#)

August 2023

- Introduced 3 new predefined post-launch actions:

- [Verify tags](#)
- [Auto Scaling group setting](#)
- Enable Refactor Spaces

[Learn more about predefined post-launch actions.](#)

- Service launch in the Israel (Tel Aviv) region.

July 2023

- You can use MGN with workloads that require FedRAMP Moderate categorization level in the AWS US East/West regions.

June 2023

- Service launch in the following regions: Europe (Zurich), Europe (Spain), Asia Pacific (Hyderabad), Asia Pacific (Melbourne).
- Introduced Import and export from local disk. You can now import and export your source servers, applications, and waves from and to a CSV file on your local disk. [Learn more about the import and export feature.](#)
- Introduced 4 new predefined post-launch actions:
 - [Configure Time Sync](#)
 - [Validate disk space](#)
 - [Verify HTTP/HTTPS response](#)
 - [Enable Amazon Inspector](#)

[Learn more about predefined post-launch actions.](#)

- Introduced global view, that allows you to manage migrations across multiple accounts using an integration with AWS Organizations. This feature provides visibility and the ability to perform actions on source servers, apps, and waves in different AWS accounts from a single console. [Learn more about global view.](#)

- Add new actions to the source server data replication process. You can now stop and start, pause and resume data replication, from the console. You can also install the AWS Replication Agent without immediately starting the data replication. [Learn more about data replication actions.](#)
- [AWS managed policy updates](#) – Updated the AWSApplicationMigrationServiceRolePolicy policy to support the global view feature.

May 2023

- Service launch in the following regions: AWS GovCloud (US-East) and AWS GovCloud (US-West).

April 2023

- [AWS managed policy updates](#) – Updated the AWSApplicationMigrationFullAccess policy to further support automation SSM documents.

March 2023

- Introduced **Import and export**, a new feature that allows you to import and export your source servers, applications, and waves from and to a CSV file. [Learn more about the import and export feature.](#)
- Added support for the following operating systems:
 - CentOS 5.5–5.11
 - RHEL 5.5–5.11
- Added support for migration of servers using the Oracle ASM Filter Driver.
- Introduced 8 new predefined post-launch actions:
 - [Conduct EC2 connectivity checks](#)
 - [Validate volume integrity](#)
 - [Verify process status](#)
 - [Convert MS-SQL license conversion](#)
 - [Install a CloudWatch Agent](#)
 - [Upgrade Windows](#)
 - [Create AMI from instance](#)

- [Join Directory Service domain](#)

[Learn more about predefined post-launch actions.](#)

- Introduced major UI enhancements to the post-launch action feature. [Learn more about the new post-launch actions layout.](#)
- Enhanced the source server page dashboard, adding migration metrics view of the displayed servers.
- Service launch in the following regions: Middle East (UAE).
- [AWS managed policy updates](#) – Updated the AWSApplicationMigrationFullAccess policy, the AWSApplicationMigrationSSMAccess policy, and the AWSApplicationMigrationReadOnlyAccess policy.

January 2023

- [AWS managed policy updates](#) – Updated the AWSApplicationMigrationEC2Access policy.

November 2022

- Introduced support for Application management. [Learn more about Applications.](#)
- Introduced support for Wave management. [Learn more about Waves.](#)
- Added support for additional launch template options. [Learn more.](#)
- Added support for post-launch custom actions. [Learn more.](#)
- Added support for no rescan upon reboot for specific operating systems. [Learn more about the no-rescan feature.](#)
- The service onboarding process has been simplified. All initial templates: replication template, launch template, and post-launch template are initialized with defaults. The templates can be modified from the Settings page. [Learn more.](#)
- Added support for SUSE 11 operating system.
- [AWS managed policy updates](#) – added one new policy and updated two existing policies. For details see [AWS MGN updates for AWS managed policies](#).

August 2022

- Added support for migration using [AWS Local Zones](#).
- Service launch in the following region: Asia Pacific (Jakarta).
- Added additional instance families to the right-sizing mechanism. [Learn more](#).

July 2022

- Added support for automatically tagging migrated resources with the required MAP program tags. [Learn more about automatic tagging](#).

June 2022

- Added support for updating vCenter credentials for agentless replication. [Learn more about updating vCenter credentials for agentless replication](#).
- Support for agent installation using temporary credentials. [Learn more about agent installation using temporary credentials](#).

May 2022

- Added support for post-launch settings. Post-launch settings allow you to control and automate actions performed after the server has been launched in AWS. [Learn more about post-launch settings](#).
- Added support for Linux SUSE SLES 12 service packs 1 and 2 operating system.

February 2022

- Added support for the following operating systems:
 - Microsoft Windows Server 2003
 - Microsoft Windows Server 2008
 - Microsoft Windows Server 2022
 - Microsoft Windows 10
- Added support for gp3 and io2 EBS volume types for replication servers.

- Added support for UEFI boot for Windows.

January 2022

- Added support for Kernel 5.15.

December 2021

- Added support for Kernels 5.8-5.14.

November 2021

- Service launch in the following regions: Europe (Paris), Europe (Milan), Middle East (Bahrain), and Africa (Cape Town).
- MGN now supports an additional replication method that does not require agent installation on each source server. This option is available for source servers running on VMware vCenter versions 6.7 and 7.0. [Learn more about agentless replication.](#)

October 2021

- Service launch in the following regions: Asia Pacific (Mumbai), Asia Pacific (Seoul), Asia Pacific (Hong Kong), Europe (London).

July 2021

- Service launch in the following regions: US West (N. California), South America (São Paulo), Canada (Central), Asia Pacific (Osaka).

April 2021

- Service initial launch in: US East (N. Virginia), US East (Ohio), US West (Oregon), Europe (Ireland), Europe (Frankfurt), Europe (Stockholm), Asia Pacific (Sydney), Asia Pacific (Singapore).

Document history for user guide

The following are the latest major policy updates for AWS Transform MGN. We also update the documentation frequently to address the feedback that you send us.

For additional details regarding new features and major updates, [see the MGN release notes](#).

Change	Description	Date
Added Amazon EBS volume settings to the default EC2 launch template	You can now configure the Volume initialization rate and Delete on termination Amazon EBS volume settings in the default EC2 launch template. For more information, see Default EC2 launch template settings .	September 8, 2026
Added AlmaLinux 8 support	Added AlmaLinux 8.6, 8.7, 8.8, 8.9, and 8.10 to the list of supported Linux operating systems.	September 6, 2026
FSx for ONTAP target storage type is now generally available	FSx for ONTAP as a target storage type for migrated workloads is now generally available . You can migrate block storage workloads directly to FSx for ONTAP as part of the same migration wave. For setup instructions, see FSx for ONTAP configuration .	August 30, 2026

Change	Description	Date
Added RHEL, Rocky Linux, and AlmaLinux 10.2 support	Added RHEL 10.2, Rocky Linux 10.2, and AlmaLinux 10.2 to the list of supported Linux operating systems.	August 30, 2026
Added Oracle Linux 9.6 support	Added Oracle Linux 9.6 to the list of supported Linux operating systems.	August 23, 2026
Added Amazon Linux 2023 with kernel 6.18 support	Added Amazon Linux 2023 with kernel 6.18 to the list of supported Linux operating systems.	August 5, 2026
Added RHEL, Rocky Linux, and AlmaLinux 9.8 support	Added RHEL 9.8, Rocky Linux 9.8, and AlmaLinux 9.8 to the list of supported Linux operating systems.	July 12, 2026
Added Windows 11 to the no-rescan upon reboot supported OS list	Added Windows 11 to the no-rescan upon reboot supported OS list.	June 18, 2026
Updated no-rescan supported OS list	Added Windows Server 2025, CentOS Stream 9, and CentOS Stream 10 to the no-rescan upon reboot supported OS list.	June 17, 2026
Added Oracle Linux 10.1 support	Added Oracle Linux 10.1 to the list of supported Linux operating systems.	June 16, 2026

Change	Description	Date
Added FSx for ONTAP as a target storage type (Public Preview)	You can now choose FSx for ONTAP as the target storage type for migrated workloads (Public Preview). Added FSx for ONTAP configuration and Troubleshooting FSx for ONTAP issues .	June 16, 2026
Updated support matrix with new OS	Added Windows 11 to the list of supported Windows operating systems.	June 10, 2026
Updated support matrix with new OS	Added Oracle Linux 9.7 and CentOS Stream 10 to the list of supported Linux operating systems.	June 4, 2026
Updated support matrix with new OS	Added CentOS Stream 9 to the list of supported Linux operating systems.	May 6, 2026
Updated network requirements	Added Amazon Linux 2023 package repository Amazon S3 bucket to the firewall allowlist and Amazon S3 VPC endpoint policy requirements for staging area subnet connectivity.	May 6, 2026
Updated supported operating systems content.	Deprecation notices updated with new entries.	December 14, 2025

Change	Description	Date
Added IPv6 support.	Added IPv6 support	December 09, 2025
Created AWS managed policy	Created the AWSApplicationMigrationNetworkMigrationCustomResource policy to provide permissions for Network Migration custom resource.	December 2, 2025
Created AWS managed policy	Created the AWSApplicationMigrationNetworkMigrationMultiAccount policy to provide permissions to automate VMware to AWS network infrastructure migration through CloudFormation.	November 10, 2025
Updated AWS managed policy	Updated multiple topics to reflect MGN adding support for IPv6	October 30, 2025
Updated AWS managed policy	Updated the AWSApplicationMigrationSSMAccess and AWSApplicationMigrationFullAccess policies to support changes in SSM.	July 3, 2025
Updated AWS managed policy	Updated the AWSApplicationMigrationServiceRolePolicy policy to support tagging network interfaces during RunInstances.	March 13, 2025

Change	Description	Date
Updated AWS managed policy	Updated the AWSApplicationMigrationEC2Access policy to support tagging network interfaces during RunInstances.	February 11, 2025
Updated AWS managed policies	Created new revisions of the following managed policies to support a change in authentication with EBS APIs: • AWSApplicationMigrationServiceRolePolicy • AWSApplicationMigrationEC2Access	January 08, 2025
Updated AWS managed policy	Updated the AWSApplicationMigrationFullAccess policy to support SecureString parameter type in SSM Parameters Store for post-migration framework actions.	March 10, 2024
Updated AWS managed policy	Updated the AWSApplicationMigrationServiceEc2InstancePolicy to support MGN to GovCloud and added SID to statements in the managed policy.	December 28, 2023

Change	Description	Date
Created AWS managed policy	Created the AWSApplicationMigrationServiceEc2InstancePolicy .	August 21, 2023
Updated AWS managed policy	Updated the AWSApplicationMigrationServiceRolePolicy policy to support the global view feature.	June 4, 2023
Updated AWS managed policy	Updated the AWSApplicationMigrationFullAccess policy to support specific automation SSM documents.	April 1, 2023
Updated AWS managed policy	Updated the AWSApplicationMigrationFullAccess policy to support both command and automation SSM documents for post-migration framework actions.	March 21, 2023
Updated AWS managed policy	Updated the AWSApplicationMigrationSSMAccess policy to support both command and automation SSM documents for the custom actions feature.	March 21, 2023

Change	Description	Date
Updated AWS managed policy	Updated the AWSApplicationMigrationReadOnlyAccess policy to support the new import and export feature.	March 21, 2023
Updated AWS managed policy	Updated the AWSApplicationMigrationEC2Access policy to support: DescribeSnapshots, DescribeImages, DescribeVolumes.	January 29, 2023