aws

User Guide

Amazon Verified Permissions

Amazon Verified Permissions: User Guide

Table of Contents

What is Amazon Verified Permissions?

Amazon Verified Permissions is a scalable, fine-grained permissions management and authorization service for custom applications built by you. Verified Permissions enables your developers to build secure applications faster by externalizing authorization and centralizing policy management and administration. Verified Permissions uses the Cedar policy language to define fine-grained permissions to protect your application's resources.

For guidance and examples for setting up a policy decision point (PDP) using Verified Permissions, see Implementing a PDP by using Amazon Verified Permissions in *AWS Prescriptive Guidance*.

Topics

Authorization in Verified Permissions

Verified Permissions provides *authorization* by verifying whether a principal is allowed to perform an action on a resource in a given context in your application. Verified Permissions presumes that the principal has been previously identified and authenticated through other means, such as by using protocols like OpenID Connect, a hosted provider like Amazon Cognito, or another authentication solution. Verified Permissions is agnostic to where the principal is managed and how they were authenticated.

Verified Permissions is a service that enables customers to create, maintain, and test policies in the AWS Management Console, programmatically using the Verified Permissions APIs, or through infrastructure as code solutions like CloudFormation. Permissions are expressed using the Cedar policy language. The client application calls authorization APIs to evaluate the Cedar policies stored with the service and provide an access decision for whether an action is permitted.

Cedar policy language

Authorization policies in Verified Permissions are written by using the Cedar policy language. Cedar is an open source language for writing authorization policies and making authorization decisions based on those policies. When you create an application, you need to ensure that only authorized principals, human users or machines, can access the application, and can do only what they're authorized to do. Using Cedar, you can decouple your business logic from the authorization logic. In your application's code, you preface requests made to your operations with a call to the Cedar authorization engine, asking "Is this request authorized?". Then, the application can either perform the requested operation if the decision is "allow", or return an error message if the decision is "deny".

Verified Permissions currently uses **Cedar version 4.7**.

For more information about Cedar, see the following:

- Cedar policy language Reference Guide
- Cedar GitHub repository

Benefits of Verified Permissions

Accelerate application development

Accelerate application development by decoupling authorization from business logic.

Verified Permissions provides integrations with popular development frameworks, making it easier to implement authorization in your applications with minimal code changes. These integrations allow you to focus on your core business logic while Verified Permissions handles the authorization decisions.

- **Express.js** – A middleware-based integration that enables you to protect API endpoints in your Express applications without modifying existing route handlers. For more information, see the section called "Using Express".

More secure applications

Verified Permissions enables developers to build more secure applications.

End-user features

Verified Permissions allows you to deliver richer end-user features for permissions management.

Related services

- **Amazon Cognito** – Amazon Cognito is an identity platform for web and mobile apps. It's a user directory, an authentication server, and an authorization service for OAuth 2.0 access tokens and AWS credentials. When you create a policy store, you have the option to build your principals and groups from an Amazon Cognito user pool. For more information, see the Amazon Cognito Developer Guide.
- **Amazon API Gateway** – Amazon API Gateway is an AWS service for creating, publishing, maintaining, monitoring, and securing REST, HTTP, and WebSocket APIs at any scale. When you create a policy store, you have the option to build your actions and resources from an API in API Gateway. For more information about API Gateway, see the API Gateway Developer Guide.
- **AWS IAM Identity Center** – With IAM Identity Center, you can manage sign-in security for your workforce identities, also known as workforce users. IAM Identity Center provides one place where you can create or connect workforce users and centrally manage their access across all their AWS accounts and applications. For more information, see the AWS IAM Identity Center User Guide.

Accessing Verified Permissions

You can work with Amazon Verified Permissions in any of the following ways.

AWS Management Console

The console is a browser-based interface to manage Verified Permissions and AWS resources. For more information about accessing Verified Permissions through the console, see How to sign in to AWS in the *AWS Sign-In User Guide*.

- Amazon Verified Permissions console

AWS Command Line Tools

You can use the AWS command line tools to issue commands at your system's command line to perform Verified Permissions and AWS tasks. Using the command line can be faster and more convenient than the console. The command line tools are also useful if you want to build scripts that perform AWS tasks.

AWS provides two sets of command line tools: the AWS Command Line Interface (AWS CLI) and the AWS Tools for Windows PowerShell. For information about installing and using the AWS CLI, see the AWS Command Line Interface User Guide. For information about installing and using the Tools for Windows PowerShell, see the AWS Tools for PowerShell User Guide.

- verifiedpermissions in the AWS CLI Command Reference
- Amazon Verified Permissions in AWS Tools for Windows PowerShell

AWS SDKs

AWS provides SDKs (software development kits) that consist of libraries and sample code for various programming languages and platforms (Java, Python, Ruby, .NET, iOS, Android, etc.). The SDKs provide a convenient way to create programmatic access to Verified Permissions and AWS. For example, the SDKs take care of tasks such as cryptographically signing requests, managing errors, and retrying requests automatically.

To learn more and download AWS SDKs, see Tools for Amazon Web Services.

The following are links to documentation for Verified Permissions resources in various AWS SDKs.

- AWS SDK for .NET
- AWS SDK for C++
- AWS SDK for Go
- AWS SDK for Java
- AWS SDK for JavaScript
- AWS SDK for PHP
- AWS SDK for Python (Boto)
- AWS SDK for Ruby
- AWS SDK for Rust

AWS CDK constructs

The AWS Cloud Development Kit (AWS CDK) is an open-source software development framework for defining cloud infrastructure in code and provisioning it through CloudFormation. Constructs, or reusable cloud components, can be used to create CloudFormation templates. These templates can then be used to deploy your cloud infrastructure.

To learn more and download AWS CDK, see AWS Cloud Development Kit.

The following are links to documentation for Verified Permissions AWS CDK resources, such as constructs.

- Amazon Verified Permissions L2 CDK Construct

Verified Permissions API

You can access Verified Permissions and AWS programmatically by using the Verified Permissions API, which lets you issue HTTPS requests directly to the service. When you use the API, you must include code to digitally sign requests using your credentials.

- Amazon Verified Permissions API Reference Guide

Pricing for Verified Permissions

Verified Permissions provides tiered pricing based on the amount of authorization requests per month made by your applications to Verified Permissions. There is also pricing for policy management actions based on the amount of cURL (client URL) policy API requests per month made by your applications to Verified Permissions.

For a complete list of charges and prices for Verified Permissions see Amazon Verified Permissions pricing.

To see your bill, go to the **Billing and Cost Management Dashboard** in the AWS Billing and Cost Management console. Your bill contains links to usage reports that provide details about your bill. To learn more about AWS account billing, see the AWS Billing User Guide.

If you have questions concerning AWS billing, accounts, and events, contact Support.

Create your first Amazon Verified Permissions policy store

For this tutorial, let's assume you're the developer of a photo sharing application and you are looking for a way to control what actions the users of the application can perform. You want to control who can add, delete, or view photos and photo albums. You also want to control what actions a user can take on their account. Can they manage their account, how about the account of a friend? To control these actions you would create policies that permit or forbid these actions based on the identity of the user. Verified Permissions offers policy stores, or containers, to house these policies.

In this tutorial we'll walk through creating a sample policy store using the Amazon Verified Permissions console. The console offers a few sample policy store options and we're going to create a **PhotoFlash** policy store. This policy store allows *principals*, such as users, to perform *actions*, such as sharing, on *resources*, such as photos or albums.

The following diagram illustrates the relationships between a principal, `User::alice`, and the actions she can take on various resources, namely her PhotoFlash account, the `VactionPhoto94.jpg` file, the photo album `alice-favorites-album`, and the user group `alice-friend-group`.

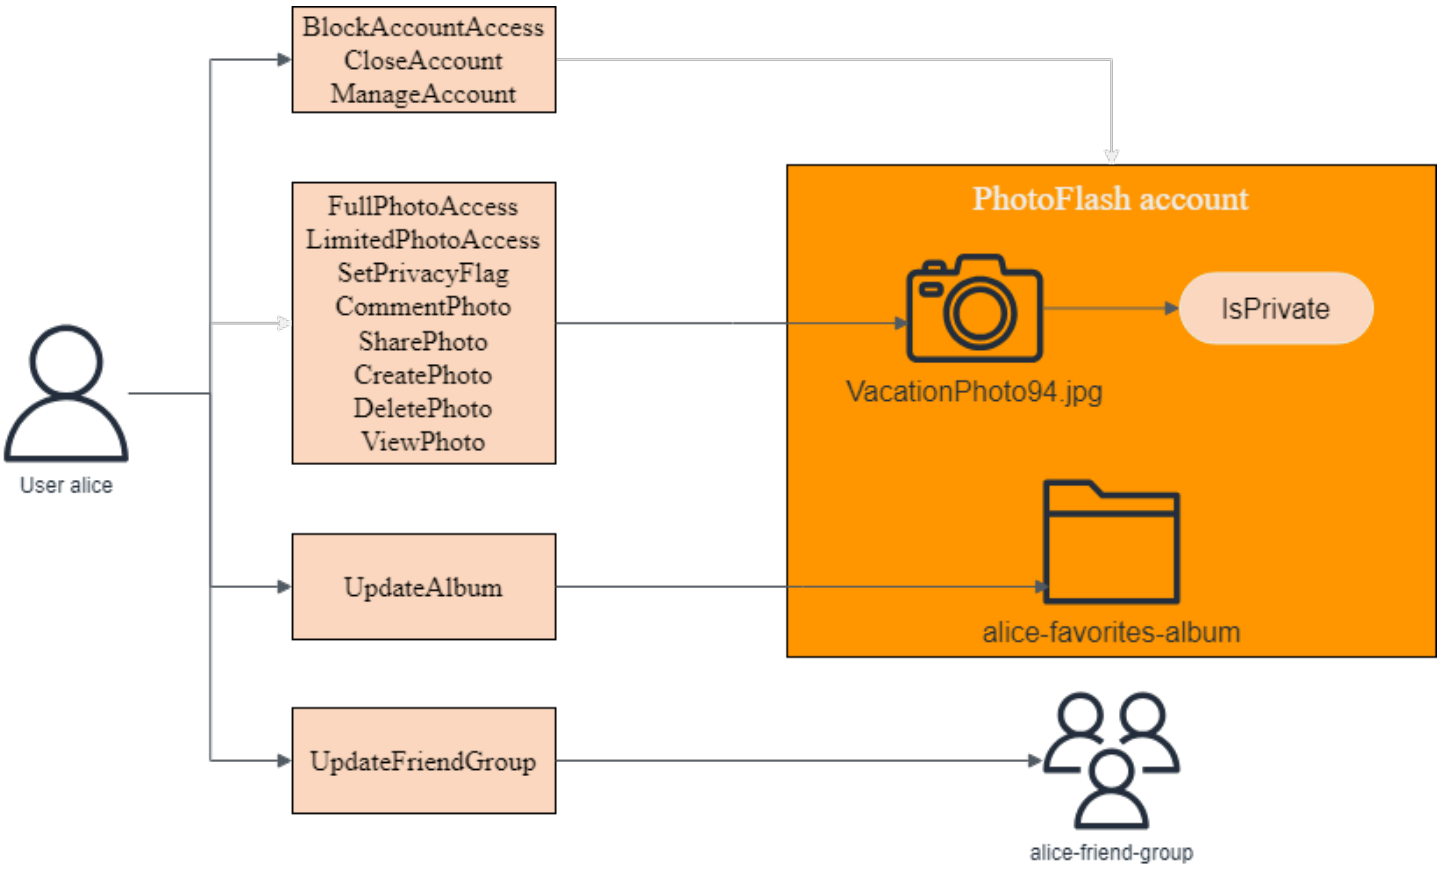

Now that you have an understanding of the **PhotoFlash** policy store, let's create the policy store and explore it.

Prerequisites

Sign up for an AWS account

To get started with AWS, you need an AWS account. For information about creating an AWS account, see Getting started with an AWS account in the *AWS Account Management Reference Guide*.

Step 1: Create a PhotoFlash policy store

In the following procedure you'll create a **PhotoFlash** policy store using the AWS console.

To create a PhotoFlash policy store

1. In the Verified Permissions console, choose **Create new policy store**.

2. For **Starting options**, choose **Start from a sample policy store**.
3. For **Sample project**, choose **PhotoFlash**.
4. Choose **Create policy store**.

Once you see the message "Created and configured policy store," choose **Go to overview** to explore your policy store.

Step 2: Create a policy

When you created the policy store, a default policy was created that allows users to have full control over their own accounts. This is a useful policy, but for our purposes, let's create a more restrictive policy to explore the nuances of Verified Permissions. If you remember the diagram we looked at earlier in the tutorial, we had a principal, `User::alice`, who could perform an action, `UpdateAlbum`, on a resource, `alice-favorites-album`. Let's add the policy that will allow Alice, and only Alice, to manage this album.

To create a policy

1. In the Verified Permissions console, choose the policy store you created in step 1.
2. In the navigation, choose **Policies**.
3. Choose **Create policy** and then choose **Create static policy**.
4. For **Policy effect**, choose **Permit**.
5. For **Principals scope**, choose **Specific principal**, then for **Specify entity type**, choose **PhotoFlash::User**, and for **Specify entity identifier**, enter **`alice`**.
6. For **Resources scope**, choose **Specific resource**, then for **Specify entity type**, choose **PhotoFlash::Album**, and for **Specify entity identifier**, enter **`alice-favorites-album`**.
7. For **Actions scope**, choose **Specific set of actions**, then for **Action(s) this policy should apply to**, select **UpdateAlbum**.
8. Choose **Next**.
9. Under **Details**, for **Policy description - optional** enter **`Policy allowing alice to update alice-favorites-album.`**.
10. Choose Create policy

Now that you've created a policy you can test it in the Verified Permissions console.

Step 3: Testing a policy store

After creating your policy store and policy, you can test them by running a simulated authorization request using the Verified Permissions test bench.

To test policy store policies

1. Open the Verified Permissions console. Choose your policy store.
2. In the navigation pane on the left, choose **Test bench**.
3. Choose **Visual mode**.
4. For **Principal**, do the following:
    a. For **Principal taking action** choose **PhotoFlash::User** and for **Specify entity identifier**, enter **`alice`**.
    b. Under **Attributes**, for **Account: Entity**, make sure that the **PhotoFlash::Account** entity is selected, and for **Specify entity identifier**, enter **`alice-account`**.
5. Under **Resource**, for **Resource that principal is acting on**, choose the **PhotoFlash::Album** resource type and for **Specify entity identifier**, enter **`alice-favorites-album`**.
6. For **Action**, choose **PhotoFlash::Action::"UpdateAlbum"** from the list of valid actions.
7. At the top of the page, choose **Run authorization request** to simulate the authorization request for the Cedar policies in the sample policy store. The test bench should display **Decision: Allow** indicating our policy is working as expected.

The following table provides additional values for the principal, resource, and action you can test with the Verified Permissions test bench. The table includes the authorization request decision based on the static policies included with the PhotoFlash sample policy store and the policy you created in step 2.

Principal value	Principal Account: Entity value	Resource value	Resource parent value	Action	Authorization decision
PhotoFlash::User | bob	**PhotoFlash::Account** | alice-account	**PhotoFlash::Album**	N/A	**PhotoFlash::Action**	Deny

Principal value	Principal Account: Entity value	Resource value	Resource parent value	Action	Authorization decision
		| alice-favorites-album		**::"UpdateAlbum"**	
PhotoFlash::User | alice	**PhotoFlash::Account |** alice-account	**PhotoFlash::Photo |** photo.jpeg	**PhotoFlash::Account |** bob-account	**PhotoFlash::Action::"ViewPhoto"**	Deny
PhotoFlash::User | alice	**PhotoFlash::Account |** alice-account	**PhotoFlash::Photo |** photo.jpeg	**PhotoFlash::Account |** alice-account	**PhotoFlash::Action::"ViewPhoto"**	Allow
PhotoFlash::User | alice	**PhotoFlash::Account |** alice-account	**PhotoFlash::Photo |** bob-photo.jpeg	**PhotoFlash::Album |** Bob-Vacation-Album	**PhotoFlash::Action::"DeletePhoto"**	Deny

Step 4: Clean up resources

After you have finished exploring your policy store, delete it.

To delete a policy store

1. In the Verified Permissions console, choose the policy store you created in step 1.
2. In the navigation, choose **Settings**.
3. Under **Delete policy store**, choose **Delete this policy store**.
4. In the **Delete this policy store?** dialog box, enter *delete*, and then choose **Delete**.

Best practices for designing an authorization model

As you prepare to use the Amazon Verified Permissions service within a software application, it can be challenging to leap immediately into writing policy statements as a first step. This would be similar to beginning development of other portions of an application by writing SQL statements or API specifications before fully deciding what the application should do. Instead, you should begin with a user experience. Then, work backwards from that experience to arrive at an implementation approach.

As you do this work, you'll find yourself asking questions such as:

- What are my resources? How are they organized? For example, do files reside within a folder?
- Does the organization of the resources play a part in the permissions model?
- What actions can principals perform on each resource?
- How do principals acquire those permissions?
- Do you want your end-users to choose from predefined permissions such as “Admin”, “Operator”, or “ReadOnly”, or should they create ad-hoc policy statements? Or both?
- Are roles global or scoped? For example, is an "operator" limited within a single tenant, or does "operator" means operator across the whole application?
- What types of queries are necessary to render the user experience? For example, do you need to list all of the resources that a principal can access to render that user's home page?
- Can users accidentally lock themselves out of their own resources? Does that need to be avoided?

The end result of this exercise is referred to as an **authorization model**; it defines the principals, resources, actions, and how they interrelate to each other. Producing this model doesn't require unique knowledge of Cedar or the Verified Permissions service. Instead, it is first and foremost a user experience design exercise, much like any other, and can manifest in artifacts such as interface mockups, logical diagrams, and an overall description of how permissions influence what users can do in the product. Cedar is designed to be flexible enough to meet customers at a model, rather than forcing the model to bend unnaturally to comply with a Cedar's implementation. As a result, gaining a crisp understanding of the desired user experience is the best way to arrive at an optimal model.

To help answer the questions and come to an optimal model, do the following:

- Review Cedar design patterns in the Cedar policy language Reference Guide.
- Consider the best practices in the Cedar policy language Reference Guide.
- Consider the best practices included on this page.

Best practices

- There isn't a canonical “correct” model
- Return 403 forbidden errors rather than 404 not found errors
- Focus on your resources beyond API operations
- Multi-tenancy considerations

There isn't a canonical “correct” model

When you design an authorization model, there is no single, uniquely correct answer. Different applications can effectively use different authorization models for similar concepts, and this is OK. For example, consider the representation of a computer's file system. When you create a file in a Unix-like operating system, it doesn't automatically inherit permissions from the parent folder. In contrast, in many other operating systems and most online file-sharing services, files do inherit permissions from its parent folder. Both choices are valid depending upon the circumstances the application is optimizing for.

The correctness of an authorization solution isn’t absolute, but should be viewed in terms of how it delivers the experience that your customers want, and whether it protects their resources in the way they expect. If your authorization model delivers on this, then it is successful.

This is why beginning your design with the desired user experience is the most helpful prerequisite to the creation of an effective authorization model.

Return 403 forbidden errors rather than 404 not found errors

It's best to return a *403 Forbidden* error to requests that include an entity, especially a resource, that doesn't correspond to any policy rather than a *404 Not found* error. This provides the highest level of security because you're not exposing whether an entity exists or not, just that the request didn't meet the policy conditions in any policy in the policy store.

Focus on your resources beyond API operations

In most applications, permissions are modeled around the resources supported. For example, a file-sharing application might represent permissions as actions that can be performed on a file or a folder. This is a good, simple model that abstracts away the underlying implementation and the backend API operations.

In contrast, other types of applications, particularly web services, frequently design permissions around the API operations themselves. For example, if a web service provides an API named `createThing()`, the authorization model might define a corresponding permission, or an `action` in Cedar named `createThing`. This works in many situations and makes it easy to understand the permissions. To invoke the `createThing` operation, you need the `createThing` action permission. Seems simple, right?

You'll find that the getting started process in the Verified Permissions console includes the option to build your resources and actions directly from an API. This is a useful baseline: a direct mapping between your policy store and the API that it authorizes for.

However, as you further develop your model, this API-focused approach may not be a good fit for applications with very granular authorization models because APIs are merely a proxy for what your customers are truly trying to protect: the underlying data and resources. If multiple APIs control access to the same resources, it can be difficult for administrators to reason about the paths to those resources and manage access accordingly.

For example, consider a user directory that contains the members of an organization. Users can be organized into groups, and one of the security goals is to prohibit discovery of group memberships by unauthorized parties. The service managing this user directory provides two API operations:

- `listMembersOfGroup`
- `listGroupMembershipsForUser`

Customers can use either of these operations to discover group membership. Therefore, the permissions administrator must remember to coordinate access to *both* operations. This is complicated further if you later choose to add a new API operation to address additional use cases, such as the following.

- `isUserInGroups` *(a new API to quickly test if a user belongs in one or more groups)*

From a security perspective, this API opens a third path for discovering group memberships, disrupting the carefully crafted permissions of the administrator.

We recommend that you focus on the underlying data and resources and their association operations. Applying this approach to the group membership example would lead to an abstract permission, such as `viewGroupMembership`, which each of the three API operations must consult.

API Name	Permissions	
`listMembersOfGroup`	requires `viewGroupMembership`	permission on the group
`listGroupMembershipsForUser`	requires `viewGroupMembership`	permission on the user
`isUserInGroups`	requires `viewGroupMembership`	permission on the user

By defining this one permission, the administrator successfully controls access to discovering group memberships, now and forever. As a tradeoff, each API operation must now document the possibly several permissions that it requires, and the administrator must consult this documentation when crafting permissions. This can be a valid tradeoff when necessary to meet your security requirements.

Multi-tenancy considerations

You might want to develop applications for use by multiple customers - businesses that consume your application, or *tenants* - and integrate them with Amazon Verified Permissions. Before you develop your authorization model, develop a multi-tenant strategy. You can manage the policies of your customers in *one shared policy store*, or assign each a *per-tenant policy store*. For more information, see Amazon Verified Permissions multi-tenant design considerations in *AWS Prescriptive Guidance*.

1. **One shared policy store**

   All tenants share a single policy store. The application sends all authorization requests to the shared policy store.

2. **Per-tenant policy store**

Each tenant has a dedicated policy store. The application will query different policy stores for an authorization decision, depending on the tenant that makes the request.

Neither strategy will have a large impact on your AWS bill. So how, then, should you design your approach? The following are common conditions that might contribute to your Verified Permissions multi-tenancy authorization strategy.

Tenant policies isolation

Isolation of the policies of each tenant from the others is important to protect tenant data. When each tenant has their own policy store, they each have their own isolated set of policies.

Authorization flow

You can identify a tenant making an authorization request with a policy store ID in the request, with per-tenant policy stores. With a shared policy store, all requests use the same policy store ID.

Templates and schema management

When your application has multiple policy stores, your policy templates and a policy store schema add a level of design and maintenance overhead in each policy store.

Global policies management

You might want to apply some *global* policies to every tenant. The level of overhead for management of global policies varies between shared and per-tenant policy store models.

Tenant off-boarding

Some tenants will contribute elements to your schema and policies that are specific to their case. When a tenant is no longer active with your organization and you want to remove their data, the level of effort varies with their level of isolation from other tenants.

Service resource quotas

Verified Permissions has resource and request-rate quotas that might influence your multi-tenancy decision. For more information about quotas, see Quotas for resources.

Comparing shared policy stores and per-tenant policy stores

Each consideration requires its own level of time and resource commitment in shared and per-tenant policy store models.

Consideration	Effort level in a shared policy store	Effort level in per-tenant policy stores
Tenant policies isolation	*Medium.* Must include tenant identifiers in policies and authorization requests.	*Low.* Isolation is default behavior. Tenant-specific policies are inaccessible to other tenants.
Authorization flow	*Low.* All queries target one policy store.	*Medium.* Must maintain mappings between each tenant and their policy store ID.
Templates and schema management	*Low.* Must make one schema work for all tenants.	*High.* Schemas and templates might be less complex individually, but changes require more coordination and complexity.
Global policies management	*Low.* All policies are global and can be centrally updated.	*High.* You must add global policies to each policy store in onboarding. Replicate global policy updates between many policy stores.
Tenant off-boarding	*High.* Must identify and delete only tenant-specific policies.	*Low.* Delete the policy store.
Service resource quotas	*High.* Tenants share resource quotas that affect policy stores like schema size, policy size per resource, and identity sources per policy store.	*Low.* Each tenant has dedicated resource quotas.

How to choose

Each multi-tenant application is different. Carefully compare the two approaches and their considerations before making an architectural decision.

If your application doesn't require tenant-specific policies and uses a single identity source, one shared policy store for all tenants is likely to be the most effective solution. This results in a simpler authorization flow and global policy management. Off-boarding a tenant using one shared policy store requires less effort because the application does not need to delete tenant-specific policies.

But if your application requires many tenant-specific policies, or uses multiple identity sources, per-tenant policy stores are likely to be most effective. You can control access to tenant policies with IAM policies that grant per-tenant permissions to each policy store. Off-boarding a tenant involves deleting their policy store; in a shared-policy-store environment, you must find and delete tenant-specific policies.

Amazon Verified Permissions policy stores

A policy store is a container for policies and policy templates. In each policy store, you can create a schema that is used to validate policies added to the policy store. In addition, you can turn on policy validation. If you add a policy to a policy store with policy validation enabled, the entity types, common types, and actions defined in the policy are validated against the schema and invalid policies are rejected.

Deletion protection prevents accidental deletion of a policy store. Deletion protection is enabled on all new policy stores created through the AWS Management Console. By contrast, it is disabled for all policy stores created through an API or SDK call.

We recommend creating one policy store per application, or one policy store per tenant for multi-tenant applications. You must specify a policy store when making an authorization request. You can also create policy store aliases to refer to your policy stores by friendly names. For more information, see Amazon Verified Permissions policy store aliases.

We recommend using *namespaces* to Cedar entities in your policy stores to prevent ambiguity. A namespace is a string prefix for a type, separated by a pair of colons (`::`) as a delimiter. For example `MyApplicationNamespace::exampleType`. Verified Permissions supports up to 100 namespaces per policy store. These namespaces help keep things straight when you're working with multiple similar applications. For example, in multi-tenant applications, using a namespace to append the name of the tenant to the types defined in the schema will make them distinct from their similar counterparts used by the other tenants. When looking at the logs for the authorization requests, you'll be able to easily indentify the tenant that processed the authorization request. For more information, see Namespaces in the *Cedar policy language Reference Guide*.

Topics

Creating Verified Permissions policy stores

You can create a policy store using the following methods:

- **Follow a guided setup** – You will define a resource type with valid actions and a principal type before creating your first policy.
- **Set up with API Gateway and an identity source**– Define your principal entities with users who sign in with an identity provider (IdP), and your actions and resource entities from an Amazon API Gateway API. We recommend this option if you want your application to authorize API requests with users' group membership or other attributes.
- **Start from a sample policy store** – Choose a pre-defined sample project policy store. We recommend this option if you are learning about Verified Permissions and want to view and test example policies.
- **Create an empty policy store** – You will define the schema and all access policies yourself. We recommend this option if you are already familiar with configuring a policy store.

Guided setup

To create a policy store using the Guided setup configuration method

The guided setup wizard leads you through the process of creating the first iteration of your policy store. You will create a schema for your first resource type, describe the actions that are applicable for that resource type, and the principal type for which you are granting permissions. You will then create your first policy. Once you've completed this wizard, you will be able to add to your policy store, extend the schema to describe other resource and principal types, and create additional policies and templates.

1. In the Verified Permissions console, select **Create new policy store**.
2. In the **Starting options** section, choose **Guided setup**.
3. Enter a **Policy store description**. This text can be whatever suits your organization as a friendly reference to the function of the current policy store, for example *Weather updates web application*.
4. In the **Details** section, type a **Namespace for your schema**. For more information about namespaces, see Namespaces in the Cedar documentation.
5. Choose **Next**.
6. On the **Resource type** window, type a name for your resource type. For example, `currentTemperature` could be a resource for the *Weather updates web application*.
7. (Optional) Choose **Add an attribute** to add resource attributes. Type the **Attribute name** and choose an **Attribute type** for each attribute of the resource. Choose whether each

attribute is **Required**. For example, `temperatureFormat` could be an attribute for the `currentTemperature` resource and be either Fahrenheit or Celsius. To remove an attribute that has been added for the resource type, choose **Remove** next to the attribute.

8. In the **Actions** field, type the actions to be authorized for the specified resource type. To add additional actions for the resource type, choose **Add an action**. For example, `viewTemperature` could be an action in the *Weather updates web application*. To remove an action that has been added for the resource type, choose **Remove** next to the action.
9. In the **Name of the principal type** field, type the name for a type of principal that will be using the specified actions for your resource type. By default, **User** is added to this field but can be replaced.
10. Choose **Next**.
11. On the **Principal type** window, choose the identity source for your principal type.

    - Choose **Custom** if the principal's ID and attributes will be provided directly by your Verified Permissions application. Choose **Add an attribute** to add principal attributes. Verified Permissions uses the specified attribute values when verifying policies against the schema. To remove an attribute that has been added for the principal type, choose **Remove** next to the attribute.
    - Choose **Cognito User Pool** if the principal's ID and attributes will be provided from an ID or access token generated by Amazon Cognito. Choose **Connect user pool**. Select the **AWS Region** and type **User pool ID** of the Amazon Cognito user pool to connect to. Choose **Connect**. For more information, see Authorization with Amazon Verified Permissions in the *Amazon Cognito Developer Guide*.
    - Choose **External OIDC provider** if the principal's ID and attributes will be extracted from an ID and/or Access token, generated by an external OIDC provider and add the provider and token details.
12. Choose **Next**.
13. In the **Policy details** section, type an optional **Policy description** for your first Cedar policy.
14. In the **Principals scope** field, choose the principals that will be granted permissions from the policy.

    - Choose **Specific principal** to apply the policy to a specific principal. Choose the principal in the **Principal that will be permitted to take actions** field and type an entity identifier for the principal. For example, `user-id` could be an entity identifier in the *Weather updates web application*.

> **Note**
>
> If you are using Amazon Cognito, the entity identifier must be formatted as `<userpool-id>|<sub>`.

- Choose **All principals** to apply the policy to all principals in your policy store.

15. In the **Resources scope** field, choose which resources that the specified principals will be authorized to act on.

    - Choose **Specific resource** to apply the policy to a specific resource. Choose the resource in the **Resource this policy should apply to** field and type an entity identifier for the resource. For example, `temperature-id` could be an entity identifier in the *Weather updates web application*.
    - Choose **All resources** to apply the policy to all resources in your policy store.

16. In the **Actions scope** field, choose which actions that the specified principals will be authorized to perform.

    - Choose **Specific set of actions** to apply the policy to specific actions. Select the check boxes next to the actions in the **Action(s) this policy should apply to** field.
    - Choose **All actions** to apply the policy to all actions in your policy store.

17. Review the policy in the **Policy preview** section. Choose **Create policy store**.

Set up with API Gateway and an identity source

To create a policy store using the Set up with API Gateway and an identity source configuration method

The API Gateway option secures APIs with Verified Permissions policies that are designed to make authorization decisions from users' groups, or *roles*. This option builds a policy store for testing authorization with identity-source groups and an API with a Lambda authorizer.

The users and their groups in an IdP become either your principals (ID tokens) or your context (access tokens). The methods and paths in an API Gateway API become the actions that your policies authorize. Your application becomes the resource. As a result of this workflow, Verified Permissions creates a policy store, a Lambda function, and an API Lambda authorizer. You must assign the Lambda authorizer to your API after you finish this workflow.

1. In the Verified Permissions console, select **Create new policy store**.
2. In the **Starting options** section, choose **Set up with API Gateway and an identity source** and select **Next**.
3. In the **Import resources and actions** step, under **API**, choose an API that will function as the model to your policy store resources and actions.

    a. Choose a **Deployment stage** from the stages configured in your API and select **Import API**. For more information about API stages, see Setting up a stage for a REST API in the Amazon API Gateway Developer Guide.

    b. Preview your **Map of imported resources and actions**.

    c. To update resources or actions, modify your API paths or methods in the API Gateway console and select **Import API** to see the updates.

    d. When you are satisfied with your choices, choose **Next**.
4. In **Identity source**, choose an **Identity provider type**. You can choose an Amazon Cognito user pool or an OpenID Connect (OIDC) IdP type.
5. If you chose **Amazon Cognito**:

    a. Choose a user pool in the same AWS Region and AWS account as your policy store.

    b. Choose the **Token type to pass to API** that you want to submit for authorization. Either token types contains user groups, the foundation of this API-linked authorization model.

    c. Under **App client validation**, you can limit the scope of a policy store to a subset of the Amazon Cognito app clients in a multi-tenant user pool. To require that user authenticate with one or more specified app clients in your user pool, select **Only accept tokens with expected app client IDs**. To accept any user who authenticates with the user pool, select **Don't validate app client IDs**.

    d. Choose **Next**.
6. If you chose **External OIDC provider:**

    a. In **Issuer URL**, enter the URL of your OIDC issuer. This is the service endpoint that provides the authorization server, signing keys, and other information about your provider, for example `https://auth.example.com`. Your issuer URL must host an OIDC discovery document at `/.well-known/openid-configuration`.

b. In **Token type**, choose the type of OIDC JWT that you want your application to submit for authorization. For more information, see Mapping Amazon Cognito tokens to schema and Mapping OIDC tokens to schema.

c. (optional) In **Token claims - optional**, choose **Add a token claim**, enter a name for the token, and select a value type.

d. In **User and group token claims**, do the following:

i. Enter a **User claim name in token** for the identity source. This is a claim, typically `sub`, from your ID or access token that holds the unique identifier for the entity to be evaluated. Identities from the connected OIDC IdP will be mapped to the user type in your policy store.

ii. Enter a **Group claim name in token** for the identity source. This is a claim, typically `groups`, from your ID or access token that contains a list of the user's groups. Your policy store will authorize requests based on the group membership.

e. In **Audience validation**, choose `Add value` and add a value that you want your policy store to accept in authorization requests.

f. Choose **Next**.

7. If you chose **Amazon Cognito**, Verified Permissions queries your user pool for groups. For OIDC providers, enter group names manually. The **Assign actions to groups** step creates policies for your policy store that permit group members to perform actions.

a. Choose or add the groups that you want to include in your policies.

b. Assign actions to each of the groups that you selected.

c. Choose **Next**.

8. In **Deploy app integration**, choose whether you want to manually attach the Lambda authorizer manually later or if you want Verified Permissions to do it for you now and review the steps that Verified Permissions will take to create your policy store and Lambda authorizer.

9. When you're ready to create the new resources, choose **Create policy store**.

10. Keep the **Policy store status** step open in your browser to monitor the progress of resource creation by Verified Permissions.

11. After some time, typically about an hour, or when the **Deploy Lambda authorizer** step shows **Success**, if you chose to attach the authorizer manually, configure your authorizer.

Verified Permissions will have created a Lambda function and a Lambda authorizer in your API. Choose **Open API** to navigate to your API.

To learn how to assign a Lambda authorizer, see Use API Gateway Lambda authorizers in the *Amazon API Gateway Developer Guide*.

a. Navigate to **Authorizers** for your API and note the name of the authorizer that Verified Permissions created.

b. Navigate to **Resources** and select a top-level method in your API.

c. Select **Edit** under **Method request settings**.

d. Set the **Authorizer** to be the authorizer name you noted earlier.

e. Expand **HTTP request headers**, enter a **Name** or `AUTHORIZATION`, and select **Required**.

f. Deploy the API stage.

g. **Save** your changes.

12. Test your authorizer with a user pool token of the **Token type** that you selected in the **Choose identity source** step. For more information about user pool sign-in and retrieving tokens, see User pool authentication flow in the *Amazon Cognito Developer Guide*.
13. Test authentication again with a user pool token in the `AUTHORIZATION` header of a request to your API.
14. Examine your new policy store. Add and refine policies.

Sample policy store

To create a policy store using the Sample policy store configuration method

1. In the **Starting options** section, choose **Sample policy store**.
2. In the **Sample project** section, choose the type of sample Verified Permissions application to use.

   - **PhotoFlash** is a sample customer-facing web application that enables users to share individual photos and albums with friends. Users can set fine-grained permissions on who is allowed to view, comment on, and re-share their photos. Account owners can also create groups of friends and organize photos into albums.

- **DigitalPetStore** is a sample application where anyone can register and become a customer. Customers can add pets for sale, search pets, and place orders. Customers who have added a pet are recorded as the pet owner. Pet owners can update the pet's details, upload a pet image, or delete the pet listing. Customers who have placed an order are recorded as the order owner. Order owners can get details on the order or cancel it. Pet store managers have administrative access.

  > **Note**
  >
  > The **DigitalPetStore** sample policy store does not include policy templates. The **PhotoFlash** and **TinyTodo** sample policy stores include policy templates.

- **TinyTodo** is a sample application that enables users to create taks and task lists. List owners can manage and share their lists and specify who can view or edit their lists.

3. A namespace for the schema of your sample policy store is automatically generated based on the sample project you chose.
4. Choose **Create policy store**.

   Your policy store is created with policies and a schema for the sample policy store you chose. For more information on template-linked policies you can create for the sample policy stores, see Amazon Verified Permissions example template-linked policies.

Empty policy store

To create a policy store using the Empty policy store configuration method

1. In the **Starting options** section, choose **Empty policy store**.
2. Choose **Create policy store**.

An empty policy store is created without a schema, which means policies are not validated. For more information about updating the schema for your policy store, see Amazon Verified Permissions policy store schema.

For more information about creating policies for your policy store, see Creating Amazon Verified Permissions static policies and Creating Amazon Verified Permissions template-linked policies.

AWS CLI

To create an empty policy store by using the AWS CLI.

You can create a policy store by using the `create-policy-store` operation.

> **Note**
>
> A policy store that you create by using the AWS CLI is empty.
>
> - To add schema, see Amazon Verified Permissions policy store schema.
> - To add policies, see Creating Amazon Verified Permissions static policies.
> - To add policy templates, see Creating Amazon Verified Permissions policy templates.

```
$ aws verifiedpermissions create-policy-store \
    --validation-settings "mode=STRICT"
{
    "arn": "arn:aws:verifiedpermissions::123456789012:policy-store/
PSEXAMPLEabcdefg111111",
    "createdDate": "2023-05-16T17:41:29.103459+00:00",
    "lastUpdatedDate": "2023-05-16T17:41:29.103459+00:00",
    "policyStoreId": "PSEXAMPLEabcdefg111111"
}
```

AWS SDKs

You can create a policy store using the `CreatePolicyStore` API. For more information, see CreatePolicyStore in the Amazon Verified Permissions API Reference Guide.

Implementing Amazon Verified Permissions in Rust with the AWS SDK

This topic provides a practical example of implementing Amazon Verified Permissions in Rust with the AWS SDK. This example shows how to develop an authorization model that can test whether a user is able to view a photo. The sample code uses the aws-sdk-verifiedpermissions crate from the AWS SDK for Rust, which offers a robust set of tools for interacting with AWS services.

Prerequisites

Before starting, ensure that you have the AWS CLI configured on your system and that you're familiar with Rust.

- For instructions on installing the AWS CLI, see AWS CLI installation guide.
- For instructions on configuring the AWS CLI, see Configuring settings for the AWS CLI and Configuration and credential file settings in the AWS CLI.
- For more information on Rust, see rust-lang.org and the AWS SDK for Rust Developer Guide.

With your environment prepared, let's explore how to implement Verified Permissions in Rust.

Test the sample code

The sample code does the following:

- Sets up the SDK client to communicate with AWS
- Creates a policy store
- Defines the structure of the policy store by adding a schema
- Adds a policy to check authorization requests
- Sends a test authorization request to verify everything is set up correctly

To test the sample code

1. Create a Rust project.
2. Replace any existing code in `main.rs` with the following code:

```
use std::time::Duration;
use std::thread::sleep;
use aws_config::BehaviorVersion;
use aws_sdk_verifiedpermissions::Client;
use aws_sdk_verifiedpermissions::{
    operation::{
        create_policy::CreatePolicyOutput,
        create_policy_store::CreatePolicyStoreOutput,
        is_authorized::IsAuthorizedOutput,
        put_schema::PutSchemaOutput,
    },
```

```
    types::{
        ActionIdentifier, EntityIdentifier, PolicyDefinition, SchemaDefinition,
 StaticPolicyDefinition, ValidationSettings
    },
};

//Function that creates a policy store in the client that's passed
async fn create_policy_store(client: &Client, valid_settings: &ValidationSettings)-
> CreatePolicyStoreOutput {
    let policy_store =
 client.create_policy_store().validation_settings(valid_settings.clone()).send().await;
    return policy_store.unwrap();
}

//Function that adds a schema to the policy store in the client
async fn put_schema(client: &Client, ps_id: &str, schema: &str) -> PutSchemaOutput
 {
    let schema =
 client.put_schema().definition(SchemaDefinition::CedarJson(schema.to_string())).policy_sto
    return schema.unwrap();
}

//Function that creates a policy in the policy store in the client
async fn create_policy(client: &Client, ps_id: &str,
 policy_definition:&PolicyDefinition) -> CreatePolicyOutput {
    let create_policy =
 client.create_policy().definition(policy_definition.clone()).policy_store_id(ps_id).send()
    return create_policy.unwrap();
}

//Function that tests the authorization request to the policy store in the client
async fn authorize(client: &Client, ps_id: &str, principal: &EntityIdentifier,
 action: &ActionIdentifier, resource: &EntityIdentifier) -> IsAuthorizedOutput {
    let is_auth =
 client.is_authorized().principal(principal.to_owned()).action(action.to_owned()).resource(
    return is_auth.unwrap();
}

#[::tokio::main]
async fn main() -> Result<(), aws_sdk_verifiedpermissions::Error> {

//Set up SDK client
    let config = aws_config::load_defaults(BehaviorVersion::latest()).await;
    let client = aws_sdk_verifiedpermissions::Client::new(&config);
```

```
//Create a policy store
    let valid_settings = ValidationSettings::builder()
    .mode({aws_sdk_verifiedpermissions::types::ValidationMode::Strict
    })
    .build()
    .unwrap();
    let policy_store = create_policy_store(&client, &valid_settings).await;
    println!(
    "Created Policy store with ID: {:?}",
    policy_store.policy_store_id
    );

//Add schema to policy store
    let schema= r#"{
        "PhotoFlash": {
            "actions": {
                "ViewPhoto": {
                    "appliesTo": {
                        "context": {
                            "type": "Record",
                            "attributes": {}
                        },
                        "principalTypes": [
                            "User"
                        ],
                        "resourceTypes": [
                            "Photo"
                        ]
                    },
                    "memberOf": []
                }
            },
            "entityTypes": {
                "Photo": {
                    "memberOfTypes": [],
                    "shape": {
                        "type": "Record",
                        "attributes": {
                            "IsPrivate": {
                                "type": "Boolean"
                            }
                        }
                    }
```

```
                },
                "User": {
                    "memberOfTypes": [],
                    "shape": {
                        "attributes": {},
                        "type": "Record"
                    }
                }
            }
        }
    }"#;
    let put_schema = put_schema(&client, &policy_store.policy_store_id,
 schema).await;
    println!(
        "Created Schema with Namespace: {:?}",
        put_schema.namespaces
    );

//Create policy
    let policy_text = r#"
        permit (
            principal in PhotoFlash::User::"alice",
            action == PhotoFlash::Action::"ViewPhoto",
            resource == PhotoFlash::Photo::"VacationPhoto94.jpg"
        );
        "#;
    let policy_definition =
 PolicyDefinition::Static(StaticPolicyDefinition::builder().statement(policy_text).build().
    let policy = create_policy(&client, &policy_store.policy_store_id,
 &policy_definition).await;
    println!(
        "Created Policy with ID: {:?}",
        policy.policy_id
    );

//Break to make sure the resources are created before testing authorization
    sleep(Duration::new(2, 0));

//Test authorization
    let principal=
 EntityIdentifier::builder().entity_id("alice").entity_type("PhotoFlash::User").build().unw
    let action =
 ActionIdentifier::builder().action_type("PhotoFlash::Action").action_id("ViewPhoto").build
```

```
    let resource =
 EntityIdentifier::builder().entity_id("VacationPhoto94.jpg").entity_type("PhotoFlash::Phot
    let auth = authorize(&client, &policy_store.policy_store_id, &principal,
 &action, &resource).await;
    println!(
        "Decision: {:?}",
        auth.decision
        );
        println!(
        "Policy ID: {:?}",
        auth.determining_policies
        );
     Ok(())
}
```

3. Run the code by entering `cargo run` in the terminal.

If the code runs correctly, the terminal will show `Decision: Allow` followed by the policy ID of the determining policy. This means you've successfully created a policy store and tested it using the AWS SDK for Rust.

Clean up resources

After you have finished exploring your policy store, delete it.

To delete a policy store

You can delete a policy store by using the `delete-policy-store` operation, replacing *PSEXAMPLEabcdefg111111* with the policy store ID you want to delete.

```
$ aws verifiedpermissions delete-policy-store \
    --policy-store-id PSEXAMPLEabcdefg111111
```

If successful, this command produces no output.

API-linked policy stores

A common use case is to use Amazon Verified Permissions to authorize user access to APIs hosted on Amazon API Gateway. Using a wizard in the AWS console, you can create role-based access policies for users managed in Amazon Cognito, or any OIDC identity provider (IdP), and deploy an AWS Lambda Authorizer that calls Verified Permissions to evaluate these policies.

To complete the wizard, choose **Set up with API Gateway and an identity provider** when you create a new policy store and follow the steps.

An API-linked policy store is created and it provisions your authorization model and resources for authorization requests. The policy store has an identity source and a Lambda authorizer that connects API Gateway to Verified Permissions. Once the policy store is created, you can authorize API requests based on users' group memberships. For example, Verified Permissions can grant access only to users who are members of the `Directors` group.

As your application grows, you can implement fine-grained authorization with user attributes and OAuth 2.0 scopes using the Cedar policy language. For example, Verified Permissions can grant access only to users who have an `email` attribute in the domain `mycompany.co.uk`.

After you have set up the authorization model for your API, your remaining responsibility is to authenticate users and generate API requests in your application, and to maintain your policy store.

To see an demo, see Amazon Verified Permissions - Quick Start Overview and Demo on the *Amazon Web Services YouTube channel*.

Topics

- How Verified Permissions authorizes API requests
- Considerations for API-linked policy stores
- Adding attribute-based access control (ABAC)
- Moving to production with AWS CloudFormation
- Troubleshooting API-linked policy stores

> **Important**
>
> Policy stores that you create with the **Set up with API Gateway and an identity source** option in the Verified Permissions console aren't intended for immediate deployment to production. With your initial policy store, finalize your authorization model and export the policy store resources to CloudFormation. Deploy Verified Permissions to production programmatically with the AWS Cloud Development Kit (AWS CDK). For more information, see Moving to production with AWS CloudFormation.

In a policy store that's linked to an API and an identity source, your application presents a user pool token in an authorization header when it makes a request to the API. The identity source of your policy store provides token validation for Verified Permissions. The token forms the `principal` in authorization requests with the IsAuthorizedWithToken API. Verified Permissions builds policies around the group membership of your users, as presented in a groups claim in identity (ID) and access tokens, for example `cognito:groups` for user pools. Your API processes the token from your application in a Lambda authorizer and submits it to Verified Permissions for an authorization decision. When your API receives the authorization decision from the Lambda authorizer, it passes the request on to your data source or denies the request.

Components of identity source and API Gateway authorization with Verified Permissions

- An Amazon Cognito user pool or OIDC IdP that authenticates and groups users. Users' tokens populate the group membership and the principal or context that Verified Permissions evaluates in your policy store.
- An API Gateway REST API. Verified Permissions defines actions from API paths and API methods, for example `MyAPI::Action::get /photo`.
- A Lambda function and a Lambda authorizer for your API. The Lambda function takes in bearer tokens from your user pool, requests authorization from Verified Permissions, and returns a decision to API Gateway. The **Set up with API Gateway and an identity source** workflow automatically creates this Lambda authorizer for you.
- A Verified Permissions policy store. The policy store identity source is your Amazon Cognito user pool or OIDC provider group. The policy store schema reflects the configuration of your API, and the policies link user groups to permitted API actions.
- An application that authenticates users with your IdP and appends tokens to API requests.

How Verified Permissions authorizes API requests

When you create a new policy store and select the **Set up with API Gateway and an identity source** option, Verified Permissions creates policy store schema and policies. The schema and policies reflect API actions and the user groups that you want to authorize to take the actions. Verified Permissions also creates the Lambda function and authorizer.

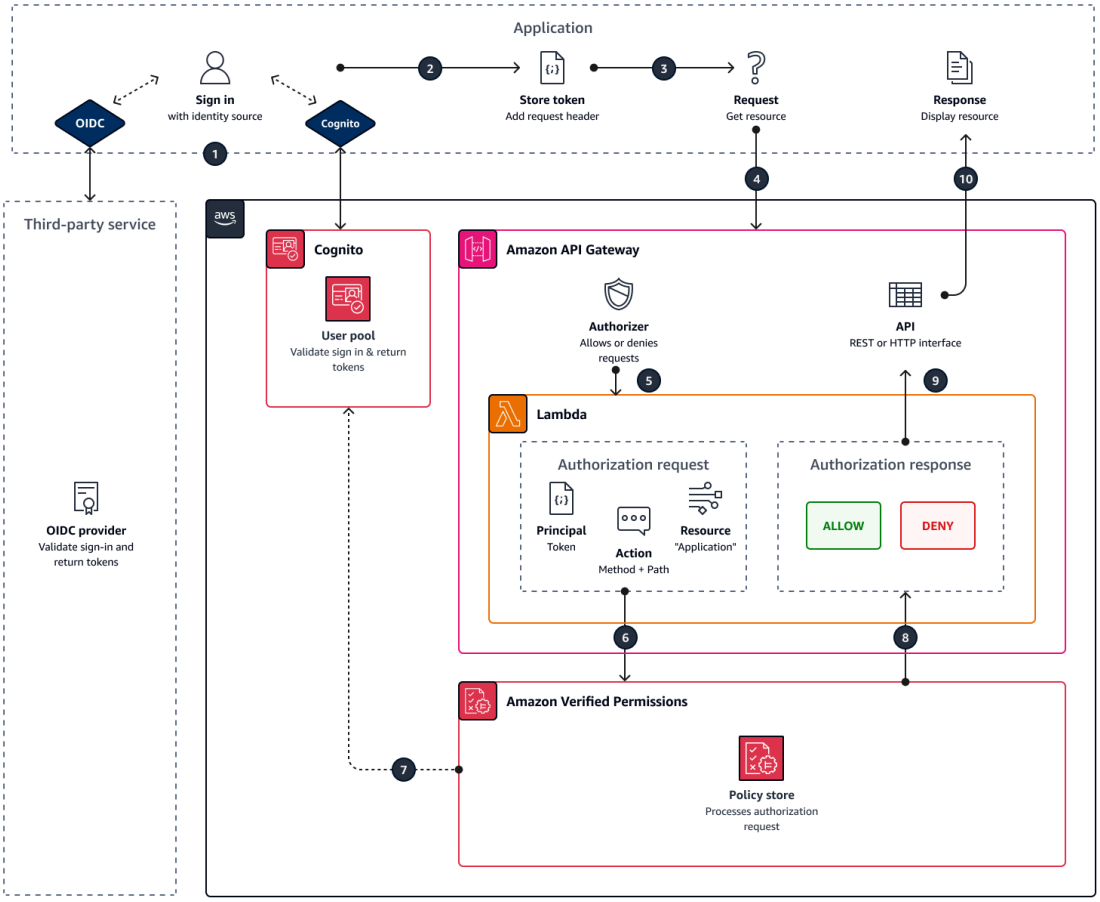

1. Your user signs in with your application through Amazon Cognito or another OIDC IdP. The IdP issues ID and access tokens with the user's information.
2. Your application stores the JWTs. For more information, see Using tokens with user pools in the *Amazon Cognito Developer Guide*..
3. Your user requests data that your application must retrieve from an external API.
4. Your application requests data from a REST API in API Gateway. It appends an ID or access token as a request header.
5. If your API has a cache for the authorization decision, it returns the previous response. If caching is disabled or the API has no current cache, API Gateway passes the request parameters to a token-based Lambda authorizer.
6. The Lambda function sends an authorization request to a Verified Permissions policy store with the IsAuthorizedWithToken API. The Lambda function passes the elements of an authorization decision:

   a. The user's token as the principal.

b. The API method combined with the API path, for example `GetPhoto`, as the action.

c. The term `Application` as the resource.

7. Verified Permissions validates the token. For more information about how Amazon Cognito tokens are validated, see Authorization with Amazon Verified Permissions in the *Amazon Cognito Developer Guide*.
8. Verified Permissions evaluates the authorization request against the policies in your policy store and returns an authorization decision.
9. The Lambda authorizer returns an `Allow` or `Deny` response to API Gateway.
10. The API returns data or an `ACCESS_DENIED` response to your application. Your application processes and displays the results of the API request.

Considerations for API-linked policy stores

When you build an API-linked policy store in the Verified Permissions console, you're creating a test for an eventual production deployment. Before you move to production, establish a fixed configuration for your API and user pool. Consider the following factors:

API Gateway caches responses

In API-linked policy stores, Verified Permissions creates a Lambda authorizer with an **Authorization caching** TTL of 120 seconds. You can adjust this value or turn off caching in your authorizer. In an authorizer with caching enabled, your authorizer returns the same response each time until the TTL expires. This can extend the effective lifetime of user pool tokens by a duration that equals the caching TTL of the requested stage.

Amazon Cognito groups can be reused

Amazon Verified Permissions determines group membership for user pool users from the `cognito:groups` claim in a user's ID or access token. The value of this claim is an array of the friendly names of the user pool groups that the user belongs to. You can't associate user pool groups with a unique identifier.

User pool groups that you delete and recreate with the same name present to your policy store as the same group. When you delete a group from a user pool, delete all references to the group from your policy store.

API-derived namespace and schema are point-in-time

Verified Permissions captures your API at a *point in time*: it only queries your API when you create your policy store. When the schema or name of your API changes, you must update your policy store and Lambda authorizer, or create a new API-linked policy store. Verified Permissions derives the policy store namespace from the name of your API.

Lambda function has no VPC configuration

The Lambda function that Verified Permissions creates for your API authorizer is launched in the default VPC. By default. APIs that have network access restricted to private VPCs can't communicate with the Lambda function that authorizes access requests with Verified Permissions.

Verified Permissions deploys authorizer resources in CloudFormation

To create an API-linked policy store, you must sign in a highly-privileged AWS principal to the Verified Permissions console. This user deploys an CloudFormation stack that creates resources across several AWS services. This principal must have the permission to add and modify resources in Verified Permissions, IAM, Lambda, and API Gateway. As a best practice, don't share these credentials with other administrators in your organization.

See Moving to production with AWS CloudFormation for an overview of the resources that Verified Permissions creates.

Adding attribute-based access control (ABAC)

A typical authentication session with an IdP returns ID and access tokens. You can pass either of these token types as a bearer token in application requests to your API. Depending on your choices when you create your policy store, Verified Permissions expects one of the two types of tokens. Both types carry information about the user's group membership. For more information about token types in Amazon Cognito, see Using tokens with user pools in the *Amazon Cognito Developer Guide*.

After you create a policy store, you can add and extend policies. For example, you can add new groups to your policies as you add them to your user pool. Because your policy store is already aware of the way that your user pool presents groups in tokens, you can permit a set of actions for any new group with a new policy.

You might also want to extend the group-based model of policy evaluation into a more precise model based on user properties. User pool tokens contain additional user information that can contribute to authorization decisions.

ID tokens

ID tokens represent a user's attributes and have a high level of fine-grained access control. To evaluate email addresses, phone numbers, or custom attributes like department and manager, evaluate the ID token.

Access tokens

Access tokens represent a user's permissions with OAuth 2.0 scopes. To add a layer of authorization or to set up requests for additional resources, evaluate the access token. For example, you can validate that a user is in the appropriate groups *and* carries a scope like `PetStore.read` that generally authorizes access to the API. User pools can add custom scopes to tokens with resource servers and with token customization at runtime.

See Mapping Amazon Cognito tokens to schema and Mapping OIDC tokens to schema for example policies that process claims in ID and access tokens.

Moving to production with AWS CloudFormation

API-linked policy stores are a way to quickly build an authorization model for an API Gateway API. They are designed to serve as a testing environment for the authorization component of your application. After you create your test policy store, spend time refining the policies, schema, and Lambda authorizer.

You might adjust the architecture of your API, requiring equivalent adjustments to your policy store schema and policies. API-linked policy stores don't automatically update their schema from API architecture–Verified Permissions only polls the API at the time you create a policy store. If your API changes sufficiently, you might have to repeat the process with a new policy store.

When your application and authorization model are ready for deployment to production, integrate the API-linked policy store that you developed with your automation processes. As a best practice, we recommend that you export the policy store schema and policies into a AWS CloudFormation template that you can deploy to other AWS accounts and AWS Regions.

The results of the API-linked policy store process are an initial policy store and a Lambda authorizer. The Lambda authorizer has several dependent resources. Verified Permissions deploys

these resources in an automatically-generated CloudFormation stack. To deploy to production, you must collect the policy store and the Lambda authorizer resources into a template. An API-linked policy store is made of the following resources:

1. AWS::VerifiedPermissions::PolicyStore: Copy your schema to the `SchemaDefinition` object. Escape " characters as \".
2. AWS::VerifiedPermissions::IdentitySource: Copy values from the output of GetIdentitySource from your test policy store and modify as needed.
3. One or more of AWS::VerifiedPermissions::Policy: Copy your policy statement to the `Definition` object. Escape " characters as \".
4. AWS::Lambda::Function, AWS::IAM::Role, AWS::IAM::Policy, AWS::ApiGateway::Authorizer, AWS::Lambda::Permission

The following template is an example policy store. You can append the Lambda authorizer resources from your existing stack to this template.

```
{
    "AWSTemplateFormatVersion": "2010-09-09",
    "Resources": {
        "MyExamplePolicyStore": {
            "Type": "AWS::VerifiedPermissions::PolicyStore",
            "Properties": {
                "ValidationSettings": {
                    "Mode": "STRICT"
                },
                "Description": "ApiGateway: PetStore/test",
                "Schema": {
                    "CedarJson": "{\"PetStore\":{\"actions\":{\"get /pets\":
{\"appliesTo\":{\"principalTypes\":[\"User\"],\"resourceTypes\":[\"Application\"],
\"context\":{\"type\":\"Record\",\"attributes\":{}}}},\"get /\":{\"appliesTo\":
{\"principalTypes\":[\"User\"],\"resourceTypes\":[\"Application\"],\"context\":{\"type
\":\"Record\",\"attributes\":{}}}},\"get /pets/{petId}\":{\"appliesTo\":{\"context
\":{\"type\":\"Record\",\"attributes\":{}},\"resourceTypes\":[\"Application\"],
\"principalTypes\":[\"User\"]}},\"post /pets\":{\"appliesTo\":{\"principalTypes\":
[\"User\"],\"resourceTypes\":[\"Application\"],\"context\":{\"type\":\"Record\",
\"attributes\":{}}}}},\"entityTypes\":{\"Application\":{\"shape\":{\"type\":\"Record\",
\"attributes\":{}}},\"User\":{\"memberOfTypes\":[\"UserGroup\"],\"shape\":{\"attributes
\":{},\"type\":\"Record\"}},\"UserGroup\":{\"shape\":{\"type\":\"Record\",\"attributes
\":{}}}}}"
                }
```

```
            }
        },
        "MyExamplePolicy": {
            "Type": "AWS::VerifiedPermissions::Policy",
            "Properties": {
                "Definition": {
                    "Static": {
                        "Description": "Policy defining permissions for testgroup
 cognito group",
                        "Statement": "permit(\nprincipal in PetStore::UserGroup::
\"us-east-1_EXAMPLE|testgroup\",\naction in [\n  PetStore::Action::\"get /\",
\n  PetStore::Action::\"post /pets\",\n  PetStore::Action::\"get /pets\",\n
 PetStore::Action::\"get /pets/{petId}\"\n],\nresource);"
                    }
                },
                "PolicyStoreId": {
                    "Ref": "MyExamplePolicyStore"
                }
            },
            "DependsOn": [
                "MyExamplePolicyStore"
            ]
        },
        "MyExampleIdentitySource": {
            "Type": "AWS::VerifiedPermissions::IdentitySource",
            "Properties": {
                "Configuration": {
                    "CognitoUserPoolConfiguration": {
                        "ClientIds": [
                            "1example23456789"
                        ],
                        "GroupConfiguration": {
                            "GroupEntityType": "PetStore::UserGroup"
                        },
                        "UserPoolArn": "arn:aws:cognito-idp:us-
east-1:123456789012:userpool/us-east-1_EXAMPLE"
                    }
                },
                "PolicyStoreId": {
                    "Ref": "MyExamplePolicyStore"
                },
                "PrincipalEntityType": "PetStore::User"
            },
            "DependsOn": [
```

```
                "MyExamplePolicyStore"
            ]
        }
    }
}
```

Troubleshooting API-linked policy stores

Use the information here to help you diagnose and fix common issues when you build Amazon Verified Permissions API-linked policy stores.

Topics

I updated my policy but the authorization decision didn't change

By default, Verified Permissions configures the Lambda authorizer to cache authorization decisions for 120 seconds. Try again after two minutes, or disable cache on your authorizer. For more information, see Enabling API caching to enhance responsiveness in the *Amazon API Gateway Developer Guide*.

I attached the Lambda authorizer to my API but it's not generating authorization requests

To begin processing requests, you must deploy the API stage that you attached your authorizer to. For more information, see Deploying a REST API in the *Amazon API Gateway Developer Guide*.

I received an unexpected authorization decision and want to review the authorization logic

The API-linked policy store process creates a Lambda function for your authorizer. Verified Permissions automatically builds the logic of your authorization decisions into the authorizer function. You can go back after you create your policy store to review and update the logic in the function.

To locate your Lambda function from the AWS CloudFormation console, choose the **Check deployment** button on the **Overview** page of your new policy store.

You can also locate your function in the AWS Lambda console. Navigate to the console in the AWS Region of your policy store and search for a function name with a prefix of `AVPAuthorizerLambda`. If you have create more than one API-linked policy store, use the **Last modified** time of your functions to correlate them with policy store creation.

I want to find logs from my Lambda authorizer

Lambda functions collect metrics and log their invocation results in Amazon CloudWatch. To review your logs, locate your function in the Lambda console and choose the **Monitor** tab. Select **View CloudWatch logs** and review the entries in the log group.

For more information about Lambda function logs, see Using Amazon CloudWatch Logs with AWS Lambda in the *AWS Lambda Developer Guide*.

My Lambda authorizer doesn't exist

After you complete setup of an API-linked policy store, you must attach the Lambda authorizer to your API. If you can't locate your authorizer in the API Gateway console, the additional resources for your policy store might have failed or not deployed yet. API-linked policy stores deploy these resources in an CloudFormation stack.

Verified Permissions displays a link with the label **Check deployment** at the end of the creation process. If you already navigated away from this screen, go to the CloudFormation console and search recent stacks for a name that's prefixed with `AVPAuthorizer-<policy store ID>`. CloudFormation provides valuable troubleshooting information in the output of a stack deployment.

For help troubleshooting CloudFormation stacks, see Troubleshooting CloudFormation in the *AWS CloudFormation User Guide*.

My API is in a private VPC and can't invoke the authorizer

Verified Permissions doesn't support access to Lambda authorizers through VPC endpoints. You must open a network path between your API and the Lambda function that serves as your authorizer.

I want to process additional user attributes in my authorization model

The API-linked policy store process derives Verified Permissions policies from the groups claim in users' tokens. To update your authorization model to consider additional user attributes, integrate those attributes in your policies.

You can map many claims in ID and access tokens from Amazon Cognito user pools to Verified Permissions policy statements. For example, most users have an `email` claim in their ID token. For more information about adding claims from your identity source to policies, see Mapping Amazon Cognito tokens to schema and Mapping OIDC tokens to schema.

I want to add new actions, action context attributes, or resource attributes

An API-linked policy store and the Lambda authorizer that it creates are a point-in-time resource. They reflect the state of your API at the time of creation. The policy store schema doesn't assign any context attributes to actions, nor any attributes or parents to the default `Application` resource.

When you add actions—paths and methods—to your API, you must update your policy store to be aware of the new actions. You must also update your Lambda authorizer to process authorization requests for the new actions. You can start again with a new policy store or you can update your existing policy store.

To update your existing policy store, locate your function. Examine the logic in the automatically-generated function and update it to process the new actions, attributes, or context. Then edit your schema to include the new actions and attributes.

Deleting policy stores

You can delete Amazon Verified Permissions policy stores using the AWS Management Console or the AWS CLI. Deleting a policy store permanently deletes the schema and any policies and policy templates in the policy store. Any policy store aliases associated with the deleted policy store will also be deleted.

Deletion protection prevents accidental deletion of a policy store. Deletion protection is enabled on all new policy stores created through the AWS Management Console. By contrast, it is disabled for all policy stores created through an API or SDK call.

You may want to delete policy stores for the following reasons:

- You have reached the quota of available policy stores in a given Region. For more information, see Quotas for resources.
- You're no longer supporting a tenant in a multi-tenant application and, therefore, no longer need that policy store.

AWS Management Console

To delete a policy store

1. Open the Verified Permissions console. Choose your policy store.
2. In the navigation pane on the left, choose **Settings**.
3. Choose **Delete this policy store**.
4. Type `delete` in the text box and choose **Delete**.

> **Note**
>
> If deletion protection is enabled, you'll need to disable it before you can choose **Delete**. To disable it, select **Disable deletion protection**.

AWS CLI

To delete a policy store

You can delete a policy store by using the `delete-policy-store` operation, replacing *PSEXAMPLEabcdefg111111* with the policy store ID you want to delete.

```
$ aws verifiedpermissions delete-policy-store \
    --policy-store-id PSEXAMPLEabcdefg111111
```

If successful, this command produces no output.

Note

If deletion protection is enabled for this policy store, you must first run the `update-policy-store` operation and disable deletion protection.

```
aws verifiedpermissions update-policy-store \
    --deletion-protection "DISABLED" \
    --policy-store-id PSEXAMPLEabcdefg111111
```

Amazon Verified Permissions policy store aliases

A policy store alias is a friendly name for a policy store. For example, policy store aliases let you refer to a policy store using `policy-store-alias/example-policy-store` instead of `PSEXAMPLEabcdefg111111`. Policy store aliases can be used in any Verified Permissions operation that accepts a `policyStoreId` input parameter.

You can create a policy store alias for a policy store by using the `CreatePolicyStoreAlias` API or by using the `AWS::VerifiedPermissions::PolicyStoreAlias` CloudFormation resource.

The Amazon Verified Permissions API provides full control of policy store aliases in each AWS account and Region. The API includes operations to create a policy store alias (`CreatePolicyStoreAlias`), view policy store alias names and policy store alias ARNs (`GetPolicyStoreAlias`, `ListPolicyStoreAliases`), and delete a policy store alias (`DeletePolicyStoreAlias`).

Topics

Properties of policy store aliases

How policy store aliases work in Amazon Verified Permissions.

A policy store alias is an independent AWS resource

A policy store alias is not a property of a policy store. The actions that you take on the policy store alias don't affect its associated policy store. You can delete the policy store alias without any effect on the associated policy store. If you delete a policy store, all policy store aliases associated with that policy store are also deleted.

Each policy store alias has an Amazon Resource Name (ARN) that uniquely identifies the policy store alias. If you specify a policy store alias as the resource in an IAM policy, the policy refers to the policy store alias, not to the associated policy store.

Each policy store alias has two formats

When you create a policy store alias, you specify the policy store alias name. Amazon Verified Permissions creates the policy store alias ARN for you.

- A policy store alias ARN is an Amazon Resource Name (ARN) that uniquely identifies the policy store alias.

```
# Alias ARN
arn:aws:verifiedpermissions:us-east-1:123456789012:policy-store-alias/example-policy-store
```

- A policy store alias name that is unique in the AWS account and Region. In the Amazon Verified Permissions API, the policy store alias name is always prefixed by `policy-store-alias/`.

```
# Alias name
policy-store-alias/example-policy-store
```

Policy store aliases are not secret

Policy store aliases may be displayed in plaintext in CloudTrail logs and other output. Do not include confidential or sensitive information in the policy store alias name.

Each policy store alias is associated with one policy store at a time

The policy store alias and its associated policy store must belong to the same AWS account and Region. You can associate a policy store alias with any policy store in the same AWS account and Region.

For example, this `ListPolicyStoreAliases` output shows that the `example-policy-store` policy store alias is associated with exactly one target policy store, which is represented by the `policyStoreId` property.

```
{
    "aliasName": "policy-store-alias/example-policy-store",
    "policyStoreId": "PSEXAMPLEabcdefg111111",
```

```
    "aliasArn": "arn:aws:verifiedpermissions:us-west-2:123456789012:policy-store-alias/
example-policy-store",
    "createdAt": "2024-01-15T12:30:00.000000+00:00",
    "state": "Active"
}
```

Multiple aliases can be associated with the same policy store

For example, you can associate the `example-policy-store` and `example-policy-store-2` aliases with the same policy store.

```
[
    {
        "aliasName": "policy-store-alias/example-policy-store",
        "policyStoreId": "PSEXAMPLEabcdefg111111",
        "aliasArn": "arn:aws:verifiedpermissions:us-west-2:123456789012:policy-store-
alias/example-policy-store",
        "createdAt": "2024-01-15T12:30:00.000000+00:00",
        "state": "Active"
    },
    {
        "aliasName": "policy-store-alias/example-policy-store-2",
        "policyStoreId": "PSEXAMPLEabcdefg111111",
        "aliasArn": "arn:aws:verifiedpermissions:us-west-2:123456789012:policy-store-
alias/example-policy-store-2",
        "createdAt": "2024-01-16T09:15:00.000000+00:00",
        "state": "Active"
    }
]
```

A policy store alias must be unique in an AWS account and Region

For example, you can have only one policy store alias with the name `example-policy-store` in each AWS account and Region. Policy store aliases are case-sensitive. You cannot change a policy store alias name. However, you can delete the policy store alias and create a new policy store alias with the desired name after the 24-hour reservation period expires.

You can create policy store aliases with the same name in different Regions. Each policy store alias will have a unique ARN. If your code refers to a policy store alias name like `policy-store-alias/example-policy-store`, you can run it in multiple Regions. In each Region, it uses a different policy store.

Policy store aliases support soft deletes and hard deletes

By default, when a policy store alias is deleted, it is soft deleted. The policy store alias name is reserved for a period of 24 hours. If you attempt to create a policy store alias with the same name during this period, the request will be rejected. During this period, `GetPolicyStoreAlias` returns the policy store alias with the `PendingDeletion` state. You can also hard delete a policy store alias by specifying `HardDelete` as the `deletionMode`. A hard-deleted policy store alias is immediately deleted and the policy store alias name becomes available for reuse immediately. For more information, see Deleting Amazon Verified Permissions policy store aliases.

You can use aliases to identify policy stores

You can use a policy store alias to identify a policy store in all operations that accept a `policyStoreId` (for example, `IsAuthorized`). In such cases, the policy store alias name must be prefixed with `policy-store-alias/`. Policy store aliases cannot be used to identify a policy store for the `DeletePolicyStore` operation.

You cannot use a policy store alias name or policy store alias ARN to identify a policy store in the `Resource` element of an IAM policy. To control access to a policy store when it is referenced through a policy store alias, see Controlling access to policy store aliases.

Creating Amazon Verified Permissions policy store aliases

You can create a policy store alias to reference a policy store using a friendly name. The name of a policy store alias must be unique per AWS account and Region. Policy store aliases may only be associated with policy stores that are owned by the same AWS account and active in the same Region as the policy store alias. Policy store aliases are separate resources with their own ARNs and IAM authorization.

By default, only 10 policy store aliases can be associated with the same policy store.

> **Note**
>
> `CreatePolicyStoreAlias` is idempotent. If you call the `CreatePolicyStoreAlias` operation with a policy store alias name and policy store ID that match an existing policy store alias, the `CreatePolicyStoreAlias` operation succeeds and returns the existing policy store alias. However, if you call the `CreatePolicyStoreAlias` operation with an existing policy store alias name but a different policy store ID, the operation returns a `ConflictException`.

AWS CLI

To create a policy store alias

You can create a policy store alias by using the CreatePolicyStoreAlias operation. The following example creates a policy store alias with the name `example-policy-store`.

```
$ aws verifiedpermissions create-policy-store-alias \
    --alias-name policy-store-alias/example-policy-store \
    --policy-store-id PSEXAMPLEabcdefg111111
{
    "aliasName": "policy-store-alias/example-policy-store",
    "policyStoreId": "PSEXAMPLEabcdefg111111",
    "aliasArn": "arn:aws:verifiedpermissions:us-west-2:123456789012:policy-store-
alias/example-policy-store",
    "createdAt": "2024-01-15T12:30:00.000000+00:00"
}
```

Retrieving Amazon Verified Permissions policy store aliases

You can retrieve information about policy store aliases using the `GetPolicyStoreAlias` operation to get details about a specific policy store alias, or the `ListPolicyStoreAliases` operation to list all policy store aliases in your AWS account and Region.

Getting a policy store alias

Use the `GetPolicyStoreAlias` operation to retrieve details about a specific policy store alias, including the associated policy store ID.

AWS CLI

To retrieve details about a policy store alias

You can retrieve a policy store alias by using the GetPolicyStoreAlias operation. The following example retrieves details for a policy store alias with the name `example-policy-store`.

```
$ aws verifiedpermissions get-policy-store-alias \
    --alias-name policy-store-alias/example-policy-store
{
    "aliasName": "policy-store-alias/example-policy-store",
```

```
    "policyStoreId": "PSEXAMPLEabcdefg111111",
    "aliasArn": "arn:aws:verifiedpermissions:us-west-2:123456789012:policy-store-
alias/example-policy-store",
    "createdAt": "2024-01-15T12:30:00.000000+00:00",
    "state": "Active"
}
```

Listing policy store aliases

Use the `ListPolicyStoreAliases` operation to list all policy store aliases in your AWS account and Region. You can use the `filter` parameter to list only policy store aliases associated with a specific policy store.

AWS CLI

To list all policy store aliases

You can list policy store aliases by using the ListPolicyStoreAliases operation. The following example lists all policy store aliases owned by the 123456789012 AWS account in the us-west-2 Region.

```
$ aws verifiedpermissions list-policy-store-aliases
{
    "policyStoreAliases": [
        {
            "aliasName": "policy-store-alias/example-policy-store",
            "policyStoreId": "PSEXAMPLEabcdefg111111",
            "aliasArn": "arn:aws:verifiedpermissions:us-west-2:123456789012:policy-
store-alias/example-policy-store",
            "createdAt": "2024-01-15T12:30:00.000000+00:00",
            "state": "Active"
        },
        {
            "aliasName": "policy-store-alias/example-policy-store-2",
            "policyStoreId": "PSEXAMPLEabcdefg111111",
            "aliasArn": "arn:aws:verifiedpermissions:us-west-2:123456789012:policy-
store-alias/example-policy-store-2",
            "createdAt": "2024-01-16T09:15:00.000000+00:00",
            "state": "Active"
        },
        {
            "aliasName": "policy-store-alias/example-policy-store-3",
```

```
            "policyStoreId": "PSEXAMPLEabcdefg222222",
            "aliasArn": "arn:aws:verifiedpermissions:us-west-2:123456789012:policy-
store-alias/example-policy-store-3",
            "createdAt": "2024-01-17T14:45:00.000000+00:00",
            "state": "Active"
        }
    ]
}
```

To list policy store aliases for a specific policy store

Use the `filter` parameter to list only aliases associated with a specific policy store.

```
$ aws verifiedpermissions list-policy-store-aliases \
    --filter '{"policyStoreId": "PSEXAMPLEabcdefg111111"}'
{
    "policyStoreAliases": [
        {
            "aliasName": "policy-store-alias/example-policy-store",
            "policyStoreId": "PSEXAMPLEabcdefg111111",
            "aliasArn": "arn:aws:verifiedpermissions:us-west-2:123456789012:policy-
store-alias/example-policy-store",
            "createdAt": "2024-01-15T12:30:00.000000+00:00",
            "state": "Active"
        },
        {
            "aliasName": "policy-store-alias/example-policy-store-2",
            "policyStoreId": "PSEXAMPLEabcdefg111111",
            "aliasArn": "arn:aws:verifiedpermissions:us-west-2:123456789012:policy-
store-alias/example-policy-store-2",
            "createdAt": "2024-01-16T09:15:00.000000+00:00",
            "state": "Active"
        }
    ]
}
```

Deleting Amazon Verified Permissions policy store aliases

You can delete a policy store alias when it is no longer needed. Deleting a policy store alias does not affect the associated policy store. Deleting a policy store deletes all policy store aliases associated with that policy store.

Amazon Verified Permissions supports two deletion modes for policy store aliases:

- **Soft delete** (default): The policy store alias enters a `PendingDeletion` state. The policy store alias name is reserved for 24 hours and cannot be reused during this period. During this period, `GetPolicyStoreAlias` returns the policy store alias with the `PendingDeletion` state. This is the default behavior when you don't specify a `deletionMode`, or when you specify `SoftDelete`.
- **Hard delete**: The policy store alias is immediately deleted. The policy store alias name becomes available for reuse immediately. To perform a hard delete, specify `HardDelete` as the `deletionMode`.

Soft deleting a policy store alias

By default, deleting a policy store alias performs a soft delete. The policy store alias enters the `PendingDeletion` state and the policy store alias name is reserved for 24 hours.

AWS CLI

To soft delete a policy store alias

You can soft delete a policy store alias by using the DeletePolicyStoreAlias operation. The following example soft deletes a policy store alias with the name `example-policy-store`.

```
$ aws verifiedpermissions delete-policy-store-alias \
    --alias-name policy-store-alias/example-policy-store
```

You can also explicitly specify the soft delete mode.

```
$ aws verifiedpermissions delete-policy-store-alias \
    --alias-name policy-store-alias/example-policy-store \
    --deletion-mode SoftDelete
```

Hard deleting a policy store alias

To immediately delete a policy store alias, specify `HardDelete` as the `deletionMode`. A hard-deleted policy store alias does not enter the `PendingDeletion` state and the policy store alias name becomes available for reuse immediately.

If you hard delete a policy store alias that was previously soft deleted, the policy store alias is immediately deleted.

> **Important**
>
> Amazon Verified Permissions is eventually consistent. If you hard delete a policy store alias and immediately recreate it to point to a different policy store, requests that reference the policy store alias may continue to resolve to the previously associated policy store for a short period of time. To avoid unexpected authorization results, allow time for the deletion to propagate before recreating a policy store alias with the same name.

AWS CLI

To hard delete a policy store alias

You can hard delete a policy store alias by using the DeletePolicyStoreAlias operation with the `deletion-mode` parameter set to `HardDelete`. The following example immediately deletes a policy store alias with the name `example-policy-store`.

```
$ aws verifiedpermissions delete-policy-store-alias \
    --alias-name policy-store-alias/example-policy-store \
    --deletion-mode HardDelete
```

Using Amazon Verified Permissions policy store aliases in API operations

Any Amazon Verified Permissions operation that accepts a `policyStoreId` parameter, such as IsAuthorized, IsAuthorizedWithToken, and GetPolicyStore, can accept a policy store alias name in place of the policy store ID.

> **Important**
>
> When you use a policy store alias as the value of a `policyStoreId` parameter, you must include the `policy-store-alias/` prefix. For example, use `policy-store-alias/example-policy-store`, not `example-policy-store`.

Using Policy store aliases in Operations

The following `IsAuthorized` command uses a policy store alias with the name `example-policy-store` to identify a policy store.

AWS CLI

```
$ aws verifiedpermissions is-authorized \
    --policy-store-id policy-store-alias/example-policy-store \
    --principal entityType=User,entityId=alice \
    --action actionType=Action,actionId=view \
    --resource entityType=Photo,entityId=photo123
```

Note

You cannot use a policy store alias in place of the `policyStoreId` field for the DeletePolicyStore operation.

Using Policy store aliases Across AWS Regions

One of the most powerful uses of aliases is in applications that run in multiple AWS Regions. For example, you might have a global application that uses different policy stores in each Region.

- In us-east-1, you want to use `PSEXAMPLEabcdefg111111`.
- In eu-west-1, you want to use `PSEXAMPLEabcdefg222222`.

You could create a different version of your application in each Region or use a dictionary or switch statement to select the right policy store for each Region. But it's much easier to create a policy store alias with the same policy store alias name in each Region. Remember that the policy store alias name is case-sensitive.

AWS CLI

```
$ aws --region us-east-1 verifiedpermissions create-policy-store-alias \
    --alias-name policy-store-alias/my-app \
    --policy-store-id PSEXAMPLEabcdefg111111
```

```
$ aws --region eu-west-1 verifiedpermissions create-policy-store-alias \
    --alias-name policy-store-alias/my-app \
    --policy-store-id PSEXAMPLEabcdefg222222
```

Then, use the policy store alias in your code. When your code runs in each Region, the policy store alias will refer to its associated policy store in that Region.

AWS CLI

```
$ aws verifiedpermissions is-authorized \
    --policy-store-id policy-store-alias/my-app \
    --principal entityType=User,entityId=alice \
    --action actionType=Action,actionId=view \
    --resource entityType=Photo,entityId=photo123
```

However, there is a risk that the policy store alias might be deleted. In that case, the application's attempts to use the policy store alias name will fail, and you might need to recreate or update the policy store alias. To mitigate this risk, be cautious about giving principals permission to manage the policy store aliases that you use in your application.

Policy store aliases are not a traffic control mechanism

A policy store alias is a stable, friendly name for a policy store. It is not a mechanism for shifting, splitting, or weighting authorization traffic between policy stores. If you are familiar with features that route traffic between versions of a resource, such as Lambda aliases, note that policy store aliases don't behave the same way. By design, a policy store alias is closer to an AWS KMS key alias: it provides a durable name for a resource, not a routing layer in front of it.

For this reason, we intentionally don't provide an `UpdatePolicyStoreAlias` operation. To change the policy store that a policy store alias points to, you delete the policy store alias and create a new policy store alias that has the same name and targets a different policy store. This is not an atomic operation, and it doesn't provide the guarantees that a traffic control mechanism would.

When you change the policy store that a policy store alias resolves to, the change doesn't take effect everywhere at the same instant:

- The updated mapping must propagate from the control plane to the data plane that evaluates authorization requests. The mapping doesn't propagate instantaneously.

- Because policy store alias resolution is eventually consistent and can be cached, a transition window exists after a change. During this window, the previously associated policy store serves some requests, and the newly associated policy store serves other requests. The length of this window is indeterminate.

During this window, your application can receive authorization decisions from either policy store. Because different policy stores can contain different policies, entities, and schema, the same request can produce different decisions depending on which policy store serves it. Policy store aliases therefore can't provide an atomic cut-over between policy stores, and they aren't a substitute for a deployment or traffic-management strategy.

> ⚠ **Don't use aliases as a cut-over mechanism**
>
> Don't build a higher-level infrastructure-as-code or deployment primitive that repoints a policy store alias in order to switch authorization traffic from one policy store to another. Because the change is not atomic, requests can be authorized against both policy stores at the same time for an indeterminate period. This can lead to inconsistent authorization decisions. To switch policy stores safely, see Performing no-downtime policy store changes.

Performing no-downtime policy store changes

Because policy store aliases don't provide an atomic cut-over (see Policy store aliases are not a traffic control mechanism), we recommend that you avoid migrating from one policy store to another by repointing a policy store alias. Instead, control the migration from within your application so that you can validate the new policy store and roll back instantly if you need to. The following approach lets you switch policy stores without downtime:

1. Create the new policy store and give it its own policy store alias, so that the current policy store and the new policy store each have a distinct, stable name. For example, keep `policy-store-alias/example-policy-store` pointing to your current policy store and create `policy-store-alias/example-policy-store-2` for the new policy store. Don't reuse or repoint a single policy store alias to perform the switch.
2. Replicate your policies, schema, and any other configuration into the new policy store, and run it in parallel with the current policy store.

3. Add a configuration value or feature flag to your application (for example, by using AppConfig) that determines which policy store alias is authoritative for authorization decisions. Don't rely on changing the alias itself to switch traffic.
4. Run the new policy store in shadow mode. Send the same `IsAuthorized` or `IsAuthorizedWithToken` request to both policy stores and compare the decisions. Record and investigate any discrepancies until the new policy store returns the decisions that you expect.
5. Use the feature flag to shift authorization decisions to the new policy store alias gradually, for example, in a phased rollout across your application hosts or user segments. Monitor authorization decisions and error rates as you proceed.
6. Use the feature flag to switch back to the original policy store alias immediately if you detect a problem. Because your application controls the switch, rollback is instant and doesn't depend on alias propagation.
7. Decommission the old policy store and its policy store alias after the new policy store is fully validated and serving all traffic.

This pattern gives your application, rather than alias propagation, control over which policy store serves each request. That control is what makes the migration safe and reversible, and it avoids the transition window that repointing a policy store alias would introduce.

Controlling access to policy store aliases

Principals who manage policy store aliases must have permission to interact with those policy store aliases and, for some operations, the policy store that the policy store alias is associated with. You can provide these permissions using IAM policies.

The following sections describe the permissions required to create and manage policy store aliases.

verifiedpermissions:CreatePolicyStoreAlias

To create a policy store alias, the principal needs the following permissions for both the policy store alias and for the associated policy store.

- `verifiedpermissions:CreatePolicyStoreAlias` for the policy store alias. Provide this permission in an IAM policy that is attached to the principal who is allowed to create the policy store alias.

The following example policy statement specifies a particular policy store alias in a `Resource` element. But you can list multiple policy store alias ARNs or specify a policy store alias pattern, such as `"sample*"`. You can also specify a `Resource` value of `"*"` to allow the principal to create any policy store alias in the AWS account and Region.

```
{
  "Sid": "IAMPolicyForCreateAlias",
  "Effect": "Allow",
  "Action": "verifiedpermissions:CreatePolicyStoreAlias",
  "Resource": "arn:aws:verifiedpermissions:us-east-1:123456789012:policy-store-alias/
example-policy-store"
}
```

- `verifiedpermissions:CreatePolicyStoreAlias` for the associated policy store. This permission must be provided in an IAM policy.

```
{
  "Sid": "PolicyStorePermissionForAlias",
  "Effect": "Allow",
  "Action": "verifiedpermissions:CreatePolicyStoreAlias",
  "Resource": "arn:aws:verifiedpermissions::123456789012:policy-store/
PSEXAMPLEabcdefg111111"
}
```

verifiedpermissions:GetPolicyStoreAlias

To get details about a specific policy store alias, the principal must have `verifiedpermissions:GetPolicyStoreAlias` permission for the policy store alias in an IAM policy.

The following example policy statement gives the principal permission to get a specific policy store alias.

```
{
  "Sid": "IAMPolicyForGetAlias",
  "Effect": "Allow",
  "Action": "verifiedpermissions:GetPolicyStoreAlias",
  "Resource": "arn:aws:verifiedpermissions:us-east-1:123456789012:policy-store-alias/
example-policy-store"
```

```
}
```

verifiedpermissions:ListPolicyStoreAliases

To list policy store aliases in the AWS account and Region, the principal must have `verifiedpermissions:ListPolicyStoreAliases` permission in an IAM policy. Because this policy is not related to any particular policy store or policy store alias resource, the value of the resource element in the policy must be `"*"`.

For example, the following IAM policy statement gives the principal permission to list all policy store aliases in the AWS account.

```
{
  "Sid": "IAMPolicyForListingAliases",
  "Effect": "Allow",
  "Action": "verifiedpermissions:ListPolicyStoreAliases",
  "Resource": "*"
}
```

verifiedpermissions:DeletePolicyStoreAlias

To delete a policy store alias, the principal needs permission for just the policy store alias.

> **Note**
>
> Deleting a policy store alias has no effect on the associated policy store, although applications that reference the policy store alias will receive errors. If you mistakenly delete a policy store alias, you can recreate it after the 24-hour reservation period.

The principal needs `verifiedpermissions:DeletePolicyStoreAlias` permission for the policy store alias. Provide this permission in an IAM policy attached to the principal who is allowed to delete the policy store alias.

The following example policy statement specifies the policy store alias in a `Resource` element. But you can list multiple policy store alias ARNs or specify a policy store alias pattern, such as `"sample*"`. You can also specify a `Resource` value of `"*"` to allow the principal to delete any policy store alias in the AWS account and Region.

```
{
  "Sid": "IAMPolicyForDeleteAlias",
  "Effect": "Allow",
  "Action": "verifiedpermissions:DeletePolicyStoreAlias",
  "Resource": "arn:aws:verifiedpermissions:us-east-1:123456789012:policy-store-alias/
example-policy-store"
}
```

Limiting Policy store alias Permissions

You can use a policy store alias to reference a policy store in any operation that accepts a `policyStoreId` field as input. When you do, Amazon Verified Permissions authorizes `verifiedpermissions:GetPolicyStoreAlias` against the policy store alias and the requested operation against the associated policy store.

For example, if the `IsAuthorized` operation is performed using a policy store alias, the principal needs both:

- `verifiedpermissions:GetPolicyStoreAlias` permission for the policy store alias
- `verifiedpermissions:IsAuthorized` permission for the associated policy store

The following example policy grants permission to call `IsAuthorized` using a specific policy store alias.

```
{
  "Sid": "IAMPolicyForAliasUsage",
  "Effect": "Allow",
  "Action": "verifiedpermissions:GetPolicyStoreAlias",
  "Resource": "arn:aws:verifiedpermissions:us-east-1:123456789012:policy-store-alias/
example-policy-store"
},
{
  "Sid": "IAMPolicyForPolicyStoreOperation",
  "Effect": "Allow",
  "Action": "verifiedpermissions:IsAuthorized",
  "Resource": "arn:aws:verifiedpermissions::123456789012:policy-store/
PSEXAMPLEabcdefg111111"
}
```

To limit which policy store aliases a principal can use, restrict the `verifiedpermissions:GetPolicyStoreAlias` permission. For example, the following policy allows the principal to use any policy store alias except those beginning with `Restricted`.

```
{
  "Sid": "IAMPolicyForAliasAllow",
  "Effect": "Allow",
  "Action": "verifiedpermissions:GetPolicyStoreAlias",
  "Resource": "arn:aws:verifiedpermissions:us-east-1:123456789012:policy-store-alias/*"
},
{
  "Sid": "IAMPolicyForAliasDeny",
  "Effect": "Deny",
  "Action": "verifiedpermissions:GetPolicyStoreAlias",
  "Resource": "arn:aws:verifiedpermissions:us-east-1:123456789012:policy-store-alias/
Restricted*"
}
```

Amazon Verified Permissions policy store schema

A *schema* is a declaration of the structure of the entity types supported by your application, and the actions your application may provide in authorization requests. To see the difference between how Verified Permissions and Cedar handles schemas, see Schema support.

For more information, see Cedar schema format in the Cedar policy language Reference Guide.

> **Note**
>
> The use of schemas in Verified Permissions is optional, but they are highly recommended for production software. When you create a new policy, Verified Permissions can use the schema to validate the entities and attributes referenced in the scope and conditions to avoid typos and mistakes in policies that can lead to confusing system behavior. If you activate policy validation, then all new policies must conform with the schema.

AWS Management Console

To create a schema

1. Open the Verified Permissions console. Choose your policy store.
2. In the navigation pane on the left, choose **Schema**.
3. Choose **Create schema**.

AWS CLI

To submit a new schema, or overwrite an existing schema by using the AWS CLI.

You can create a policy store by running a AWS CLI command similar to the following example.

Consider a schema that contains the following Cedar content:

```
{
    "MySampleNamespace": {
        "actions": {
            "remoteAccess": {
```

```
                "appliesTo": {
                    "principalTypes": [ "Employee" ]
                }
            }
        },
        "entityTypes": {
            "Employee": {
                "shape": {
                    "type": "Record",
                    "attributes": {
                        "jobLevel": {"type": "Long"},
                        "name": {"type": "String"}
                    }
                }
            }
        }
    }
}
```

You must first escape the JSON into a single line string, and preface it with a declaration of its data type: `cedarJson`. The following example uses the following contents of `schema.json` file that contains the escaped version of the JSON schema.

> **Note**
>
> The example here is line wrapped for readability. You must have the entire file on a single line for the command to accept it.

```
{"cedarJson": "{\"MySampleNamespace\": {\"actions\": {\"remoteAccess\": {\"appliesTo
\":
{\"principalTypes\": [\"Employee\"]}}},\"entityTypes\": {\"Employee\": {\"shape\":
{\"attributes\": {\"jobLevel\": {\"type\": \"Long\"},\"name\": {\"type\": \"String
\"}},
\"type\": \"Record\"}}}}}"}
```

```
$ aws verifiedpermissions put-schema \
    --definition file://schema.json \
    --policy-store PSEXAMPLEabcdefg111111
{
    "policyStoreId": "PSEXAMPLEabcdefg111111",
```

```
    "namespaces": [
        "MySampleNamespace"
    ],
    "createdDate": "2023-07-17T21:07:43.659196+00:00",
    "lastUpdatedDate": "2023-08-16T17:03:53.081839+00:00"
}
```

AWS SDKs

You can create a policy store using the `PutSchema` API. For more information, see PutSchema in the Amazon Verified Permissions API Reference Guide.

Editing policy store schemas

When you select **Schema** in the Amazon Verified Permissions console, the **Entity types** and **Actions** that make up your schema are displayed. You can view edit your schema in either **Visual mode** or **JSON mode**. Visual mode lets you update the schema by adding new types and actions using various wizards. Using JSON mode, you can start updating the JSON code of the schema directly in the JSON editor.

Visual Mode

The visual schema editor begins with a series of diagrams that illustrate the relationships between the entities in your schema. Choose **Expand** to maximize your view of the diagrams. There are two diagrams available:

- **Actions diagram** – The **Actions diagram** view lists the types of **Principals** you have configured in your policy store, the **Actions** they are eligible to perform, and the **Resources** that they are eligible to perform actions on. The lines between entities indicate your ability to create a policy that allows a principal to take an action on a resource. If your actions diagram doesn't indicate a relationship between two entities, you must create that relationship between them before you can allow or deny it in policies. Select an entity to see a properties overview and drill down to view full details. Choose **Filter by this [action | resource type | principal type]** to see an entity in a view with only its own connections.
- **Entity types diagram** – The **Entity types diagram** focuses on the relationships between principals and resources. When you want to understand the complex nested parent relationships in your schema, review this diagram. Hover over an entity to drill down into the parent relationships that it has.

Under the diagrams are list views of the **Entity types** and **Actions** in your schema. The list view is useful when you want to immediately view the details of a specific action or entity type. Select any entity to view details.

To edit a Verified Permissions schema in Visual mode

1. Open the Verified Permissions console. Choose your policy store.
2. In the navigation pane on the left, choose **Schema**.
3. Choose **Visual mode**. Review the entity-relationship diagrams and plan the changes that you want to make to your schema. You can optionally **Filter by** one entity to examine its individual connections to other entities.
4. Choose **Edit schema**.
5. In the **Details** section, type a **Namespace** for your schema.
6. In the **Entity types** section, choose **Add new entity type**.
7. Type the name of the entity.
8. (Optional) Choose **Add a parent** to add parent entities that the new entity is a member of. To remove a parent that has been added to the entity, choose **Remove** next to the name of the parent.
9. Choose **Add an attribute** to add attributes to the entity. Type the **Attribute name** and choose the **Attribute type** for each attribute of the entity. Verified Permissions uses the specified attribute values when verifying policies against the schema. Select whether each attribute is **Required**. To remove an attribute that has been added to the entity, choose **Remove** next to the attribute.
10. Choose **Add entity type** to add the entity to the schema.
11. In the **Actions** section, choose **Add new action**.
12. Type the name of the action.
13. (Optional) Choose **Add a resource** to add resource types for which the action applies to. To remove a resource type that has been added to the action, choose **Remove** next to the name of the resource type.
14. (Optional) Choose **Add a principal** to add a principal type that the action applies to. To remove a principal type that has been added to the action, choose **Remove** next to the name of the principal type.

15. Choose **Add an attribute** to add attributes that can be added to the context of an action in your authorization requests. Enter the **Attribute name** and choose the **Attribute type** for each attribute. Verified Permissions uses the specified attribute values when verifying policies against the schema. Select whether each attribute is **Required**. To remove an attribute that has been added to the action, choose **Remove** next to the attribute.
16. Choose **Add action**.
17. After all the entity types and actions have been added to the schema, choose **Save changes**.

JSON mode

While making updates, you'll notice the JSON editor validates your code against JSON syntax and will identify errors and warnings as you edit, making it easier for you to find issues quickly. In addition, you don't need to worry about the formatting of the JSON, simply choose **Format JSON** once you've made your updates and the format will be updated to match expected JSON formatting.

To edit a Verified Permissions schema in JSON mode

1. Open the Verified Permissions console. Choose your policy store.
2. In the navigation pane on the left, choose **Schema**.
3. Choose **JSON mode** and then choose **Edit schema**.
4. Enter the content of your JSON schema in the **Contents** field. You can't save updates to your schema until you resolve all syntax errors. You can choose **Format JSON** to format the JSON syntax of your schema with the recommended spacing and indentation.
5. Choose **Save changes**.

Enabling Amazon Verified Permissions policy validation mode

You can set the policy validation mode in Verified Permissions to control whether policy changes are validated against the schema in your policy store.

> **Important**
>
> When you turn on policy validation, all attempts to create or update a policy or policy template are validated against the schema in the policy store. Verified Permissions rejects the request attempt if validation fails. For this reason, we recommend leaving validation off while you're developing your application and turning it on for testing and leaving it on while your application is in production.

AWS Management Console

To set the policy validation mode for a policy store

1. Open the Verified Permissions console. Choose your policy store.
2. Choose **Settings**.
3. In the **Policy validation mode** section, choose **Modify**.
4. Do one of the following:
   - To activate policy validation and enforce that all policy changes must be validated against your schema, choose the **Strict (recommended)** radio button.
   - To turn off policy validation for policy changes, choose the **Off** radio button. Type `confirm` to confirm that updates to policies will no longer be validated against your schema.
5. Choose **Save changes**.

AWS CLI

To set the validation mode for a policy store

You can change the validation mode for a policy store by using the UpdatePolicyStore operation and specifying a different value for the ValidationSettings parameter.

```
$ aws verifiedpermissions update-policy-store \
    --validation-settings "mode=OFF",
    --policy-store-id PSEXAMPLEabcdefg111111
{
    "createdDate": "2023-05-17T18:36:10.134448+00:00",
    "lastUpdatedDate": "2023-05-17T18:36:10.134448+00:00",
    "policyStoreId": "PSEXAMPLEabcdefg111111",
    "validationSettings": {
        "Mode": "OFF"
    }
}
```

For more information, see Policy validation in the *Cedar policy language Reference Guide*.

Amazon Verified Permissions policies

A *policy* is a statement that either permits or forbids a *principal* to take one or more *actions* on a *resource*. Each policy is evaluated independently of every other policy. For more information about how Cedar policies are structured and evaluated, see Cedar policy validation against schema in the Cedar policy language Reference Guide.

You can optionally assign a policy name to a policy. Policy names must be unique for all policies within the policy store and prefixed with `name/`. You can use a policy name in place of the policy ID in control plane operations that accept a `policyId` parameter. The following example uses a policy name to retrieve a policy with `GetPolicy`.

```
$ aws verifiedpermissions get-policy \
    --policy-id name/example-policy \
    --policy-store-id PSEXAMPLEabcdefg111111
```

Important

When you write Cedar policies that reference principals, resources and actions, you can define the unique identifiers used for each of those elements. We strongly recommend that you follow these best practices:

- **Use universally unique identifiers (UUIDs) for all principal and resource identifiers.**

  For example, if user `jane` leaves the company, and you later let someone else use the name `jane`, then that new user automatically gets access to everything granted by policies that still reference `User::"jane"`. Cedar can't distinguish between the new user and the old. This applies to both principal and resource identifiers. Always use identifiers that are guaranteed unique and never reused to ensure that you don't unintentionally grant access because of the presence of an old identifier in a policy.

  Where you use a UUID for an entity, we recommend that you follow it with the // comment specifier and the 'friendly' name of your entity. This helps to make your policies easier to understand. For example: principal == Role::"a1b2c3d4-e5f6-a1b2-c3d4-EXAMPLE11111", // administrators

- **Do not include personally identifying, confidential, or sensitive information as part of the unique identifier for your principals or resources.** These identifiers are included in log entries shared in AWS CloudTrail trails.

Topics

Creating Amazon Verified Permissions static policies

You can create a static policy for principals to permit or forbid them from performing specified actions on specified resources for your application. A static policy has specific values included for the `principal` and `resource` and are ready to be used in authorization decisions.

AWS Management Console

To create a static policy

1. Open the Verified Permissions console. Choose your policy store.
2. In the navigation pane on the left, choose **Policies**.
3. Choose **Create policy** and then choose **Create static policy**.

> **Note**
>
> If you have a policy statement you'd like to use, skip to **Step 8** and paste the policy into the **Policy** section on the next page.

4. In the **Policy effect** section, choose whether the policy will **Permit** or **Forbid** when a request matches the policy. If you choose **Permit**, the policy allows the principals to perform the actions on the resources. Conversely, if you choose **Forbid**, the policy doesn't allow the principals to perform the actions on the resources.

5. In the **Principals scope** field, choose the scope of the principals that the policy will apply to.
   - Choose **Specific principal** to apply the policy to a specific principal. Specify the entity type and identifier for the principal that will be permitted or forbidden to take the actions specified in the policy.
   - Choose **Group of principals** to apply the policy to a group of principals. Type the principal group name in the **Group of principals** field.
   - Choose **All principals** to apply the policy to all principals in your policy store.
6. In the **Resources scope** field, choose the scope of the resources that the policy will apply to.
   - Choose **Specific resources** to apply the policy to a specific resource. Specify the entity type and identifier for the resource that the policy should apply to.
   - Choose **Group of resources** to apply the policy to a group of resources. Type the resource group name in the **Group of resources** field.
   - Choose **All resources** to apply the policy to all resources in your policy store.
7. In the **Actions scope** section, choose the scope of the resources that the policy will apply to.
   - Choose **Specific set of actions** to apply the policy to a set of actions. Select the check boxes next to the actions to apply the policy.
   - Choose **All actions** to apply the policy to all actions in your policy store.
8. Choose **Next**.
9. In the **Policy** section, review your Cedar policy. You can choose **Format** to format the syntax of your policy with the recommended spacing and indentation. For more information, see Basic policy construction in Cedar in the Cedar policy language Reference Guide.
10. In the **Details** section, type an optional description of the policy.
11. Choose **Create policy**.

AWS CLI

To create a static policy

You can create a static policy by using the CreatePolicy operation. The following example creates a simple static policy.

```
$ aws verifiedpermissions create-policy \
    --definition "{ \"static\": { \"Description\": \"MyTestPolicy\", \"Statement\":
 \"permit(principal,action,resource) when {principal.owner == resource.owner};\"}}"
 \
    --policy-store-id PSEXAMPLEabcdefg111111
{
"Arn": "arn:aws:verifiedpermissions::123456789012:policy/PSEXAMPLEabcdefg111111/
SPEXAMPLEabcdefg111111",
    "createdDate": "2023-05-16T20:33:01.730817+00:00",
    "lastUpdatedDate": "2023-05-16T20:33:01.730817+00:00",
    "policyId": "SPEXAMPLEabcdefg111111",
    "policyStoreId": "PSEXAMPLEabcdefg111111",
    "policyType": "STATIC"
}
```

To create a policy with a policy name

You can optionally specify a policy name when creating a policy. The name must be unique for all policies within the policy store and prefixed with `name/`. You can use the name in place of the policy ID.

```
$ aws verifiedpermissions create-policy \
    --definition "{ \"static\": { \"Statement\": \"permit(principal, action,
 resource in Album::\\\"public_folder\\\");\"}}" \
    --policy-store-id PSEXAMPLEabcdefg111111 \
    --name name/example-policy
{
    "createdDate": "2023-06-12T20:33:37.382907+00:00",
    "lastUpdatedDate": "2023-06-12T20:33:37.382907+00:00",
    "policyId": "SPEXAMPLEabcdefg111111",
    "policyStoreId": "PSEXAMPLEabcdefg111111",
    "policyType": "STATIC",
    "resource": {
        "entityId": "public_folder",
        "entityType": "Album"
    }
}
```

> **Note**
>
> If you specify a name that is already associated with another policy in the policy store, you receive a `ConflictException` error.

Editing Amazon Verified Permissions static policies

You can edit an existing static policy in your policy store. You can only directly update static policies. To change a template-linked policy, you must update the policy template. For more information, see Editing Amazon Verified Permissions policy templates.

You can change the following elements of a static policy:

- The `action` referenced by the policy.
- A condition clause, such as `when` and `unless`.

You can't change the following elements of a static policy. To change any of these elements you will need to delete and re-created the policy.

- A policy from a static policy to a template-linked policy.
- The effect of a static policy from `permit` or `forbid`.
- The `principal` referenced by a static policy.
- The `resource` referenced by a static policy.

AWS Management Console

To edit a static policy

1. Open the Verified Permissions console. Choose your policy store.
2. In the navigation pane on the left, choose **Policies**.
3. Choose the radio button next to the static policy to edit and then choose **Edit**.
4. In the **Policy body** section, update the `action` or condition clause of your static policy. You can't update the policy effect, `principal`, or `resource` of the policy.
5. Choose **Update policy**.

> **Note**
>
> If policy validation is enabled in the policy store, then updating a static policy causes Verified Permissions to validate the policy against the schema in the policy store. If the updated static policy doesn't pass validation, the operation fails and the update isn't saved.

AWS CLI

To edit a static policy

You can edit a static policy by using the UpdatePolicy operation. The following example edits a simple static policy.

The example uses the file `definition.txt` to contain the policy definition.

```
{
    "static": {
        "description":  "Grant everyone of janeFriends UserGroup access to the
 vacationFolder Album",
        "statement": "permit(principal in UserGroup::\"janeFriends\", action,
 resource in Album::\"vacationFolder\" );"
    }
}
```

The following command references that file.

```
$ aws verifiedpermissions create-policy \
    --definition file://definition.txt \
    --policy-store-id PSEXAMPLEabcdefg111111

{
    "createdDate": "2023-06-12T20:33:37.382907+00:00",
    "lastUpdatedDate": "2023-06-12T20:33:37.382907+00:00",
    "policyId": "SPEXAMPLEabcdefg111111",
    "policyStoreId": "PSEXAMPLEabcdefg111111",
    "policyType": "STATIC",
    "principal": {
        "entityId": "janeFriends",
```

```
        "entityType": "UserGroup"
    },
    "resource": {
        "entityId": "vacationFolder",
        "entityType": "Album"
    }
}
```

To update the name of a policy

You can set or update a policy name when updating a policy. The name must be unique for all policies within the policy store and prefixed with `name/`. If you don't include the name field in the update request, the existing name is unchanged. To remove a name, set it to an empty string.

```
$ aws verifiedpermissions update-policy \
    --policy-id SPEXAMPLEabcdefg111111 \
    --policy-store-id PSEXAMPLEabcdefg111111 \
    --definition file://definition.txt \
    --name name/example-policy
{
    "createdDate": "2023-06-12T20:33:37.382907+00:00",
    "lastUpdatedDate": "2023-06-12T20:47:42.804511+00:00",
    "policyId": "SPEXAMPLEabcdefg111111",
    "policyStoreId": "PSEXAMPLEabcdefg111111",
    "policyType": "STATIC",
    "principal": {
        "entityId": "janeFriends",
        "entityType": "UserGroup"
    },
    "resource": {
        "entityId": "vacationFolder",
        "entityType": "Album"
    }
}
```

Adding context

Context is the information that's relevant to policy decisions, but not part of the identity of your principal, action, or resource. Access token claim are context. You might want to allow an action only from a set of source IP addresses, or only if your user has signed in with MFA. Your

application has access to this contextual session data and must populate it to authorization requests. The context data in a Verified Permissions authorization request must be JSON-formatted in a `contextMap` element.

The examples that illustrate this content come from a sample policy store. To follow along, create the **DigitalPetStore** sample policy store in your testing environment.

The following context object declares one of each Cedar data type for an application based on the sample **DigitalPetStore** policy store.

```
"context": {
  "contextMap": {
    "AccountCodes": {
      "set": [
        {
          "long": 111122223333
        },
        {
          "long": 444455556666
        },
        {
          "long": 123456789012
        }
      ]
    },
    "approvedBy": {
    "entityIdentifier": {
      "entityId": "Bob",
      "entityType": "DigitalPetStore::User"
    }
    },
    "MfaAuthorized": {
      "boolean": true
    },
    "NetworkInfo": {
      "record": {
        "IPAddress": {
          "string": "192.0.2.178"
        },
        "Country": {
          "string": "United States of America"
        },
        "SSL": {
```

```
          "boolean": true
        }
    }
    },
    "RequestedOrderCount": {
      "long": 4
    },
    "UserAgent": {
      "string": "My UserAgent 1.12"
    }
  }
}
```

Data types in authorization context

Boolean

A binary `true` or `false` value. In the example, the boolean value of `true` for `MfaAuthenticated` indicates that the customer has performed multi-factor authentication before requesting to view their order.

Set

A collection of context elements. Set members can be all the same type, like in this example, or of different types, including a nested set. In the example, the customer is associated with 3 different accounts.

String

A sequence of letters, numbers, or symbols, enclosed in " characters. In the example, the `UserAgent` string represents the browser that the customer used to request to view their order.

Long

An integer. In the example, the `RequestedOrderCount` indicates that this request is part of a batch that resulted from the customer asking to view four of their past orders.

Record

A collection of attributes. You must declare these attributes in the request context. A policy store with a schema must include this entity and the attributes of the entity in the schema. In the example, the `NetworkInfo` record contains information about the user's originating IP, the geolocation of that IP as determined by the client, and encryption in transit.

EntityIdentifier

A reference to an entity and attributes declared in the `entities` element of the request. In the example, the user's order was approved by employee Bob.

To test this example context in the example **DigitalPetStore** app, you must update your request `entities`, your policy store schema, and the static policy with the description **Customer Role - Get Order**.

Modifying DigitalPetStore to accept authorization context

Initially, **DigitalPetStore** is not a very complex policy store. It doesn't include any preconfigured policies or context attributes to support the context that we have presented. To evaluate an example authorization request with this context information, make the following modifications to your policy store and your authorization request. For context examples with access token information as the context, see Mapping Amazon Cognito access tokens and Mapping OIDC access tokens.

Schema

Apply the following updates to your policy store schema to support the new context attributes. Update `GetOrder` in `actions` as follows.

```
"GetOrder": {
  "memberOf": [],
  "appliesTo": {
    "resourceTypes": [
      "Order"
    ],
    "context": {
      "type": "Record",
      "attributes": {
        "AccountCodes": {
          "type": "Set",
          "required": true,
          "element": {
            "type": "Long"
          }
        },
        "approvedBy": {
          "name": "User",
```

```
          "required": true,
          "type": "Entity"
        },
        "MfaAuthorized": {
          "type": "Boolean",
          "required": true
        },
        "NetworkInfo": {
          "type": "NetworkInfo",
          "required": true
        },
        "RequestedOrderCount": {
          "type": "Long",
          "required": true
        },
        "UserAgent": {
          "required": true,
          "type": "String"
        }
      }
    },
    "principalTypes": [
      "User"
    ]
  }
}
```

To reference the `record` data type named `NetworkInfo` in your request context, create a commonType construct in your schema by adding the following to your schema before `actions`. A `commonType` construct is a shared set of attributes that you can apply to different entities.

```
"commonTypes": {
  "NetworkInfo": {
    "attributes": {
      "IPAddress": {
        "type": "String",
        "required": true
      },
      "SSL": {
        "required": true,
        "type": "Boolean"
      },
```

```
      "Country": {
        "required": true,
        "type": "String"
      }
    },
    "type": "Record"
  }
},
```

Policy

The following policy sets up conditions that must be fulfilled by each of the provided context elements. It builds on the existing static policy with the description **Customer Role - Get Order**. This policy initially only requires that the principal that makes a request is the owner of the resource.

```
permit (
    principal in DigitalPetStore::Role::"Customer",
    action in [DigitalPetStore::Action::"GetOrder"],
    resource
) when {
    principal == resource.owner &&
    context.AccountCodes.contains(111122223333) &&
    context.approvedBy in DigitalPetStore::Role::"Employee" &&
    context.MfaAuthorized == true &&
    context.NetworkInfo.Country like "*United States*" &&
    context.NetworkInfo.IPAddress like "192.0.2.*" &&
    context.NetworkInfo.SSL == true &&
    context.RequestedOrderCount <= 4 &&
    context.UserAgent like "*My UserAgent*"
};
```

We have now required that the request to retrieve an order meets the additional context conditions that we added to the request.

1. The user must have signed in with MFA.
2. The user's web browser `User-Agent` must contain the string `My UserAgent`.
3. The user must have requested to view 4 or fewer orders.
4. One of the user's account codes must be 111122223333.

5. The user's IP address must originate in the United States, they must be on an encrypted session, and their IP address must begin with `192.0.2.`.
6. An employee must have approved their order. In the `entities` element of the authorization request, we will declare a user Bob who has the role of `Employee`.

Request body

After you configure your policy store with the appropriate schema and policy, you can present this authorization request to the Verified Permissions API operation IsAuthorized. Note that the `entities` segment contains a definition of Bob, a user with a role of `Employee`.

```
{
  "principal": {
    "entityType": "DigitalPetStore::User",
    "entityId": "Alice"
  },
  "action": {
    "actionType": "DigitalPetStore::Action",
    "actionId": "GetOrder"
  },
  "resource": {
    "entityType": "DigitalPetStore::Order",
    "entityId": "1234"
  },
  "context": {
    "contextMap": {
      "AccountCodes": {
        "set": [
          {"long": 111122223333},
          {"long": 444455556666},
          {"long": 123456789012}
        ]
      },
      "approvedBy": {
        "entityIdentifier": {
          "entityId": "Bob",
          "entityType": "DigitalPetStore::User"
        }
      },
      "MfaAuthorized": {
        "boolean": true
      },
```

```
      "NetworkInfo": {
        "record": {
          "Country": {"string": "United States of America"},
          "IPAddress": {"string": "192.0.2.178"},
          "SSL": {"boolean": true}
        }
      },
      "RequestedOrderCount":{
        "long": 4
      },
      "UserAgent": {
        "string": "My UserAgent 1.12"
      }
    }
  },
  "entities": {
    "entityList": [
      {
        "identifier": {
          "entityType": "DigitalPetStore::User",
          "entityId": "Alice"
        },
        "attributes": {
          "memberId": {
            "string": "801b87f2-1a5c-40b3-b580-eacad506d4e6"
          }
        },
        "parents": [
          {
            "entityType": "DigitalPetStore::Role",
            "entityId": "Customer"
          }
        ]
      },
      {
        "identifier": {
          "entityType": "DigitalPetStore::User",
          "entityId": "Bob"
        },
        "attributes": {
          "memberId": {
            "string": "49d9b81e-735d-429c-989d-93bec0bcfd8b"
          }
        },
```

```
        "parents": [
          {
            "entityType": "DigitalPetStore::Role",
            "entityId": "Employee"
          }
        ]
      },
      {
        "identifier": {
          "entityType": "DigitalPetStore::Order",
          "entityId": "1234"
        },
        "attributes": {
          "owner": {
            "entityIdentifier": {
              "entityType": "DigitalPetStore::User",
              "entityId": "Alice"
            }
          }
        },
        "parents": []
      }
    ]
  },
  "policyStoreId": "PSEXAMPLEabcdefg111111"
}
```

Using the Amazon Verified Permissions test bench

Use the Verified Permissions test bench to test and troubleshoot Verified Permissions policies by running authorization requests against them. The test bench uses the parameters that you specify to determine whether the Cedar policies in your policy store would authorize the request. You can toggle between **Visual mode** and **JSON mode** while testing authorization requests. For more information about how Cedar policies are structured and evaluated, see Basic policy construction in Cedar in the Cedar policy language Reference Guide.

> **Note**
>
> When you make an authorization request using Verified Permissions, you can provide the list of principals and resources as part of the request in the **Additional entities** section.

However, you can't include the details about the actions. They must be specified in the schema or inferred from the request. You can't put an action in the **Additional entities** section.

For a visual overview and demonstration of the test bench, see Amazon Verified Permissions - Policy Creation and Testing (Primer Series #3) on the AWS YouTube channel.

Visual mode

> **Note**
>
> You must have a schema defined in your policy store to use the **Visual mode** of the test bench.

To test policies in Visual mode

1. Open the Verified Permissions console. Choose your policy store.
2. In the navigation pane on the left, choose **Test bench**.
3. Choose **Visual mode**.
4. In the **Principal** section, choose the **Principal taking action** from the principal types in your schema. Type an identifier for the principal in the text box.
5. (Optional) Choose **Add a parent** to add parent entities for the specified principal. To remove a parent that has been added to the principal, choose **Remove** next to the name of the parent.
6. Specify the **Attribute value** for each attribute of the specified principal. The test bench uses the specified attribute values in the simulated authorization request.
7. In the **Resource** section, choose the **Resource that principal is acting on**. Type an identifier for the resource in the text box.
8. (Optional) Choose **Add a parent** to add parent entities for the specified resource. To remove a parent that has been added to the resource, choose **Remove** next to the name of the parent.
9. Specify the **Attribute value** for each attribute of the specified resource. The test bench uses the specified attribute values in the simulated authorization request.

10. In the **Action** section, choose the **Action that principal is taking** from the list of valid actions for the specified principal and resource.
11. Specify the **Attribute value** for each attribute of the specified action. The test bench uses the specified attribute values in the simulated authorization request.
12. (Optional) In the **Additional entities** section, choose **Add entity** to add entities to be evaluated for the authorization decision.
13. Choose the **Entity Identifier** from the dropdown list and type the entity identifier.
14. (Optional) Choose **Add a parent** to add parent entities for the specified entity. To remove a parent that has been added to the entity, choose **Remove** next to the name of the parent.
15. Specify the **Attribute value** for each attribute of the specified entity. The test bench uses the specified attribute values in the simulated authorization request.
16. Choose **Confirm** to add the entity to the test bench.
17. Choose **Run authorization request** to simulate the authorization request for the Cedar policies in your policy store. The test bench displays the decision to allow or deny the request along with information about the policies satisfied or the errors encountered during evaluation.

JSON mode

To test policies in JSON mode

1. Open the Verified Permissions console. Choose your policy store.
2. In the navigation pane on the left, choose **Test bench**.
3. Choose **JSON mode**.
4. In the **Request details** section, if you have a schema defined, choose the **Principal taking action** from the principal types in your schema. Type an identifier for the principal in the text box.

   If you do not have a schema defined, type the principal in the **Principal taking action** text box.
5. If you have a schema defined, choose the **Resource** from the resource types in your schema. Type an identifier for the resource in the text box.

   If you do not have a schema defined, type the resource in the **Resource** text box.

6. If you have a schema defined, choose the **Action** from the list of valid actions for the specified principal and resource.

   If you do not have a schema defined, type the action in the **Action** text box.
7. Enter the context of the request to simulate in the **Context** field. The request context is additional information that can be used for authorization decisions.
8. In the **Entities** field, enter the hierarchy of the entities and their attributes to be evaluated for the authorization decision.
9. Choose **Run authorization request** to simulate the authorization request for the Cedar policies in your policy store. The test bench displays the decision to allow or deny the request along with information about the policies satisfied or the errors encountered during evaluation.

Policy size per resource

Verified Permissions limits the total size of all policies that reference the same resource. This limit is the *Policy size per resource* quota, and its default value is 200,000 bytes. When you call `CreatePolicy` or `UpdatePolicy`, Verified Permissions calculates the new total for the referenced resource. If the new total exceeds the quota, the request fails with a `ServiceQuotaExceededException`.

Policies that don't specify a resource apply to all resources. These policies share a separate total for the `"unspecified"` resource, which is limited to the same quota value.

For the list of all Verified Permissions quotas, see Quotas for Amazon Verified Permissions.

Topics

- How policy sizes are calculated
- Staying within the quota
- Template-linked policy size example

How policy sizes are calculated

Each policy contributes to the total for a resource as follows:

- **Static policies** – A static policy contributes its full size. This is the same value that is measured against the *Policy size* quota.

- **Template-linked policies** – A template-linked policy contributes the combined length of the principal and resource used to instantiate it. If you don't specify the principal or resource, its length is counted as 0. The size of the policy template body itself doesn't count toward this quota.

Each policy also adds a fixed overhead of up to a few hundred bytes to the total. Template-linked policies add more overhead than static policies.

Staying within the quota

If your application approaches the *Policy size per resource* quota, consider the following strategies:

- **Use the appropriate policy type.** If many policies grant similar permissions on the same resource, policy templates might reduce your policy store's storage consumption. Each template-linked policy stores only its principal and resource. The policy template stores the shared statement once, but each linked policy also carries a fixed storage overhead. Templates only pay for themselves when the shared statement is large enough to offset that overhead, roughly 200 bytes or more. Below that threshold, individual static policies consume less space. For more information, see Amazon Verified Permissions policy templates and template-linked policies.
- **Specify principals and resources in the policy scope.** Policies that don't specify a resource share a single total for the `"unspecified"` resource across the policy store. Scoping each policy to a specific resource spreads your policies across separate totals, one for each resource.
- **Group your principals and resources.** Instead of writing a policy for each principal, write a policy for a group and add principals to the group. Grouping resources into containers can similarly reduce the number of policies that reference a single resource.
- **Request a quota increase.** You can request an increase to this quota if your policy design meets certain constraints. To qualify, your policies must specify the principal or resource in the policy scope. An increase might also reduce the number of hierarchy parents that your principals and resources can have. Even with an increased quota, a 200,000-byte sub-limit still applies to policies that reference the same principal and the same resource. You cannot request an increase to this sub-limit.

Template-linked policy size example

To find how template-linked policies contribute to the *Policy size per resource* quota, add the length of the principal and the length of the resource. If you don't specify a resource, Verified Permissions counts its size toward the `"unspecified"` resource total.

Review the following template:

```
permit (
  principal in ?principal,
  action in
    [MyApplication::Action::"ContentManagement",
     MyApplication::Action::"AccountAdministration",
     MyApplication::Action::"BillingOperations"],
  resource in ?resource
)
when {
  resource has authorizedTenantIds &&
  principal has tenantId &&
  resource.authorizedTenantIds.contains(principal.tenantId)
};
```

Create the following policies from that template:

```
TemplateLinkedPolicy {
  policyId: "policy1",
  templateId: "template1",
  principal: User::"alice",
  resource: Photo::"car.jpg"
}

TemplateLinkedPolicy {
  policyId: "policy2",
  templateId: "template1",
  principal: User::"bob",
  resource: Photo::"boat.jpg"
}

TemplateLinkedPolicy {
  policyId: "policy3",
  templateId: "template1",
  principal: User::"jane",
```

```
  resource: Photo::"car.jpg"
}

TemplateLinkedPolicy {
  policyId: "policy4",
  templateId: "template1",
  principal: User::"jane",
  resource
}
```

Calculate the size of those policies by counting the bytes in the `principal` and `resource` for each one. Verified Permissions measures size in UTF-8 bytes. Characters in the printable ASCII range count as 1 byte each, but other characters can use more. The following examples use only ASCII characters, so each character counts as 1 byte. These calculations exclude the fixed overhead. Your actual totals will be higher.

The size of `policy1` is the length of the principal `User::"alice"` (13) plus the length of the resource `Photo::"car.jpg"` (16). The total is 13 + 16 = 29 bytes.

The size of `policy2` is the length of the principal `User::"bob"` (11) plus the length of the resource `Photo::"boat.jpg"` (17). The total is 11 + 17 = 28 bytes.

The size of `policy3` is the length of the principal `User::"jane"` (12) plus the length of the resource `Photo::"car.jpg"` (16). The total is 12 + 16 = 28 bytes.

The size of `policy4` is the length of the principal `User::"jane"` (12) plus the length of the resource (0). The total is 12 + 0 = 12 bytes.

Because `policy2` is the only policy that references the resource `Photo::"boat.jpg"`, the total resource size is 28 bytes.

Because `policy1` and `policy3` both reference the resource `Photo::"car.jpg"`, the total resource size is 29 + 28 = 57 bytes.

Because `policy4` is the only policy that references the `"unspecified"` resource, the total resource size is 12 bytes.

Amazon Verified Permissions example policies

Some of the policy examples included here are basic Cedar policy examples and some are Verified Permissions-specific. The basic ones link to the Cedar policy language Reference Guide and are

included there. For more information about Cedar policy syntax, see Basic policy construction in Cedar in the Cedar policy language Reference Guide.

Policy examples

- Allows access to individual entities
- Allows access to groups of entities
- Allows access for any entity
- Allows access for attributes of an entity (ABAC)
- Denies access
- Uses bracket notation to reference token attributes
- Uses dot notation to reference attributes
- Reflects Amazon Cognito ID token attributes
- Reflects OIDC ID token attributes
- Reflects Amazon Cognito access token attributes
- Reflects OIDC access token attributes

Uses bracket notation to reference token attributes

This following example shows how you might create a policy that uses bracket notation to reference token attributes.

For more information about using token attributes in policies in Verified Permissions, see Mapping Amazon Cognito tokens to schema and Mapping OIDC tokens to schema.

```
permit (
    principal in MyCorp::UserGroup::"us-west-2_EXAMPLE|MyUserGroup",
    action,
    resource
) when {
    principal["cognito:username"] == "alice" &&
    principal["custom:employmentStoreCode"] == "petstore-dallas" &&
    principal has email && principal.email == "alice@example.com" &&
    context["ip-address"] like "192.0.2.*"
};
```

Uses dot notation to reference attributes

This following example shows how you might create a policy that uses dot notation to reference attributes.

For more information about using token attributes in policies in Verified Permissions, see Mapping Amazon Cognito tokens to schema and Mapping OIDC tokens to schema.

```
permit(principal, action, resource)
when {
    principal.cognito.username == "alice" &&
    principal.custom.employmentStoreCode == "petstore-dallas" &&
    principal.tenant == "x11app-tenant-1" &&
    principal has email && principal.email == "alice@example.com"
};
```

Reflects Amazon Cognito ID token attributes

This following example shows how you might create a policy references ID token attributes from Amazon Cognito.

For more information about using token attributes in policies in Verified Permissions, see Mapping Amazon Cognito tokens to schema and Mapping OIDC tokens to schema.

```
permit (
    principal in MyCorp::UserGroup::"us-west-2_EXAMPLE|MyUserGroup",
    action,
    resource
) when {
    principal["cognito:username"] == "alice" &&
    principal["custom:employmentStoreCode"] == "petstore-dallas" &&
    principal.tenant == "x11app-tenant-1" &&
    principal has email && principal.email == "alice@example.com"
};
```

Reflects OIDC ID token attributes

This following example shows how you might create a policy references ID token attributes from an OIDC provider.

For more information about using token attributes in policies in Verified Permissions, see Mapping Amazon Cognito tokens to schema and Mapping OIDC tokens to schema.

```
permit (
    principal in MyCorp::UserGroup::"MyOIDCProvider|MyUserGroup",
    action,
    resource
) when {
    principal.email_verified == true && principal.email == "alice@example.com" &&
    principal.phone_number_verified == true && principal.phone_number like "+1206*"
};
```

Reflects Amazon Cognito access token attributes

This following example shows how you might create a policy references access token attributes from Amazon Cognito.

For more information about using token attributes in policies in Verified Permissions, see Mapping Amazon Cognito tokens to schema and Mapping OIDC tokens to schema.

```
permit(principal, action in [MyApplication::Action::"Read",
 MyApplication::Action::"GetStoreInventory"], resource)
when {
    context.token.client_id == "52n97d5afhfiu1c4di1k5m8f60" &&
    context.token.scope.contains("MyAPI/mydata.write")
};
```

Reflects OIDC access token attributes

This following example shows how you might create a policy references access token attributes from an OIDC provider.

For more information about using token attributes in policies in Verified Permissions, see Mapping Amazon Cognito tokens to schema and Mapping OIDC tokens to schema.

```
permit(
    principal,
    action in [MyApplication::Action::"Read",
 MyApplication::Action::"GetStoreInventory"],
    resource
)
```

```
when {
    context.token.client_id == "52n97d5afhfiu1c4di1k5m8f60" &&
    context.token.scope.contains("MyAPI-read")
};
```

Amazon Verified Permissions policy templates and template-linked policies

In Verified Permissions, policy templates are policies with placeholders for the `principal`, `resource`, or both. Policy templates alone can't be used to handle authorization requests. To handle authorization requests, a *template-linked policy* must be created based on a policy template. Policy templates allow a policy to be defined once and then used with multiple principals and resources. Updates to the policy template are reflected across all policies that use the template. For more information, see Cedar policy templates in the Cedar policy language Reference Guide.

You can optionally assign a policy template name to a policy template. Policy template names must be unique within the policy store and prefixed with `name/`. You can use a policy template name in place of the policy template ID in control plane operations that accept a `policyTemplateId` parameter. Only `GetPolicyTemplate` and `ListPolicyTemplates` return the name in the output. The following example uses a policy template name to retrieve a policy template with `GetPolicyTemplate`.

```
$ aws verifiedpermissions get-policy-template \
    --policy-template-id name/example-policy-template \
    --policy-store-id PSEXAMPLEabcdefg111111
```

For example, the following policy template provides `Read`, `Edit`, and `Comment` permissions for the principal and resource that use the policy template.

```
permit(
  principal == ?principal,
  action in [Action::"Read", Action::"Edit", Action::"Comment"],
  resource == ?resource
);
```

If you were to create a policy named `Editor` based on this template, when a principal is designated as an editor for a specific resource, your application would create a policy that provides permissions for the principal to read, edit, and comment on the resource.

Unlike static policies, template-linked policies are dynamic. Take the previous example, if you were to remove the `Comment` action from the policy template, any policy linked to, or based on, that

template would be updated accordingly and the principals specified in the policies would no longer be able to comment on the corresponding resources.

For more template-linked policy examples, see Amazon Verified Permissions example template-linked policies.

Creating Amazon Verified Permissions policy templates

You can create policy templates in Verified Permissions using the AWS Management Console, the AWS CLI, or the AWS SDKs. Policy templates allow a policy to be defined once and then used with multiple principals and resources. Once you create a policy template you can then create template-linked policies to use the policy templates with specific principals and resources. For more information, see Creating Amazon Verified Permissions template-linked policies.

AWS Management Console

To create a policy template

1. Open the Verified Permissions console. Choose your policy store.
2. In the navigation pane on the left, choose **Policy templates**.
3. Choose **Create policy template**.
4. In the **Details** section, type a **Policy template description**.
5. In the **Policy template body** section, use placeholders `?principal` and `?resource` to allow policies created based on this template to customize permissions they grant. You can choose **Format** to format the syntax of your policy template with the recommended spacing and indentation.
6. Choose **Create policy template**.

AWS CLI

To create a policy template

You can create a policy template by using the CreatePolicyTemplate operation. The following example creates a policy template with a placeholder for the principal.

The file `template1.txt` contains the following.

```
"VacationAccess"
```

```
permit(
    principal in ?principal,
    action == Action::"view",
    resource == Photo::"VacationPhoto94.jpg"
);
```

```
$ aws verifiedpermissions create-policy-template \
    --description "Template for vacation picture access"
    --statement file://template1.txt
    --policy-store-id PSEXAMPLEabcdefg111111
{
    "createdDate": "2023-05-18T21:17:47.284268+00:00",
    "lastUpdatedDate": "2023-05-18T21:17:47.284268+00:00",
    "policyStoreId": "PSEXAMPLEabcdefg111111",
    "policyTemplateId": "PTEXAMPLEabcdefg111111"
}
```

To create a policy template with a policy template name

You can optionally specify a policy template name when creating a policy template. The name must be unique for all policy templates within the policy store and prefixed with `name/`. You can use the name in place of the policy template ID.

```
$ aws verifiedpermissions create-policy-template \
    --description "Template for vacation picture access" \
    --statement file://template1.txt \
    --policy-store-id PSEXAMPLEabcdefg111111 \
    --name name/example-policy-template
{
    "createdDate": "2023-06-12T20:47:42.804511+00:00",
    "lastUpdatedDate": "2023-06-12T20:47:42.804511+00:00",
    "policyStoreId": "PSEXAMPLEabcdefg111111",
    "policyTemplateId": "PTEXAMPLEabcdefg111111"
}
```

> **Note**
>
> If you specify a name that is already associated with another policy template in the policy store, you receive a `ConflictException` error.

Creating Amazon Verified Permissions template-linked policies

You can create template-linked policies, or policies that are based on a policy template, using the AWS Management Console, AWS CLI, or the AWS SDKs. Template-linked policies stay linked to their policy templates. If you change the policy statement in the policy template, any policies linked to that template automatically use the new statement for all authorization decisions made from that moment forward.

For template-linked policy examples, see Amazon Verified Permissions example template-linked policies.

AWS Management Console

To create a template-linked policy by instantiating a policy template

1. Open the Verified Permissions console. Choose your policy store.
2. In the navigation pane on the left, choose **Policies**.
3. Choose **Create policy** and then choose **Create template-linked policy**.
4. Choose the radio button next to the policy template to use and then choose **Next**.
5. Type the **Principal** and **Resource** to be used for this specific instance of the template-linked policy. The specified values are displayed in the **Policy statement preview** field.

   > **Note**
   >
   > The **Principal** and **Resource** values must have the same formatting as static policies. For example, to specify the `AdminUsers` group for the principal, type `Group::"AdminUsers"`. If you type `AdminUsers`, a validation error is displayed.

6. Choose **Create template-linked policy**.

   The new template-linked policy is displayed under **Policies**.

AWS CLI

To create a template-linked policy by instantiating a policy template

You can create a template-linked policy that references an existing policy template and that specifies values for any placeholders used by the template.

The following example creates a template-linked policy that uses a template with the following statement:

```
permit(
    principal in ?principal,
    action == PhotoFlash::Action::"view",
    resource == PhotoFlash::Photo::"VacationPhoto94.jpg"
);
```

It also uses the following `definition.txt` file to supply the value for the `definition` parameter:

```
{
    "templateLinked": {
        "policyTemplateId": "PTEXAMPLEabcdefg111111",
        "principal": {
            "entityType": "PhotoFlash::User",
            "entityId": "alice"
        }
    }
}
```

The output shows both the resource, which it gets from the template, and the principal, which it gets from the definition parameter

```
$ aws verifiedpermissions create-policy \
    --definition file://definition.txt
    --policy-store-id PSEXAMPLEabcdefg111111
{
    "createdDate": "2023-05-22T18:57:53.298278+00:00",
    "lastUpdatedDate": "2023-05-22T18:57:53.298278+00:00",
    "policyId": "TPEXAMPLEabcdefg111111",
    "policyStoreId": "PSEXAMPLEabcdefg111111",
    "policyType": "TEMPLATELINKED",
    "principal": {
        "entityId": "alice",
        "entityType": "PhotoFlash::User"
    },
    "resource": {
        "entityId": "VacationPhoto94.jpg",
        "entityType": "PhotoFlash::Photo"
    }
```

```
}
```

Editing Amazon Verified Permissions policy templates

You can edit, or update, policy templates in Verified Permissions using the AWS Management Console, the AWS CLI, or the AWS SDKs. Editing a policy template will automatically update the policies that are linked to, or based on, the template so take care when editing the policy templates and make sure you don't accidentally introduce a change that breaks your application.

You can change the following elements of a policy template:

- The `action` referenced by the policy template
- A condition clause, such as `when` and `unless`

You can't change the following elements of a policy template. To change any of these elements you will need to delete and re-created the policy template.

- The effect of a policy template from `permit` or `forbid`
- The `principal` referenced by a policy template
- The `resource` referenced by a policy template

AWS Management Console

To edit your policy templates

1. Open the Verified Permissions console. Choose your policy store.
2. In the navigation pane on the left, choose **Policy templates**. The console displays all of the policy templates you created in the current policy store.
3. Choose the radio button next to a policy template to display details about the policy template, such as when the policy template was created, updated, and the policy template contents.
4. Choose **Edit** to edit your policy template. Update the **Policy description** and **Policy body** as necessary and then choose **Update policy template**.
5. You can delete a policy template by choosing the radio button next to a policy template and then choosing **Delete**. Choose **OK** to confirm deleting the policy template.

AWS CLI

To edit a policy template

You can create a static policy by using the UpdatePolicy operation. The following example updates the specified policy template by replacing its policy body with a new policy defined in a file.

Contents of file `template1.txt`:

```
permit(
    principal in ?principal,
    action == Action::"view",
    resource in ?resource)
when {
    principal has department && principal.department == "research"
};
```

```
$ aws verifiedpermissions update-policy-template \
    --policy-template-id PTEXAMPLEabcdefg111111 \
    --description "My updated template description" \
    --statement file://template1.txt \
    --policy-store-id PSEXAMPLEabcdefg111111
{
    "createdDate": "2023-05-17T18:58:48.795411+00:00",
    "lastUpdatedDate": "2023-05-17T19:18:48.870209+00:00",
    "policyStoreId": "PSEXAMPLEabcdefg111111",
    "policyTemplateId": "PTEXAMPLEabcdefg111111"
}
```

To update the name of a policy template

You can set or update a policy template name when updating a policy template. The name must be unique for all policy templates within the policy store and prefixed with `name/`. If you don't include the name field in the update request, the existing name is unchanged. To remove a name, set it to an empty string.

```
$ aws verifiedpermissions update-policy-template \
    --policy-template-id PTEXAMPLEabcdefg111111 \
    --statement file://template1.txt \
    --policy-store-id PSEXAMPLEabcdefg111111 \
```

```
    --name name/example-policy-template
{
    "createdDate": "2023-05-17T18:58:48.795411+00:00",
    "lastUpdatedDate": "2023-05-17T19:18:48.870209+00:00",
    "policyStoreId": "PSEXAMPLEabcdefg111111",
    "policyTemplateId": "PTEXAMPLEabcdefg111111"
}
```

Amazon Verified Permissions example template-linked policies

When you create a policy store in Verified Permissions using the **Sample policy store** method, your policy store is created with predefined policies, policy templates, and a schema for the sample project you chose. The following Verified Permissions template-linked policy examples can be used with the sample policy stores and their respective policies, policy templates, and schemas.

PhotoFlash examples

The following example shows how you might create a template-linked policy that uses the policy template **Grant limited access to non-private shared photos** with an individual user and photo.

> **Note**
>
> Cedar policy language considers an entity to be `in` itself. Therefore, `principal in User::"Alice"` is equivalent to `principal == User::"Alice"`.

```
permit (
 principal in PhotoFlash::User::"Alice",
 action in PhotoFlash::Action::"SharePhotoLimitedAccess",
 resource in PhotoFlash::Photo::"VacationPhoto94.jpg"
 );
```

The following example shows how you might create a template-linked policy that uses the policy template **Grant limited access to non-private shared photos** with an individual user and album.

```
permit (
 principal in PhotoFlash::User::"Alice",
 action in PhotoFlash::Action::"SharePhotoLimitedAccess",
 resource in PhotoFlash::Album::"Italy2023"
```

```
 );
```

The following example shows how you might create a template-linked policy that uses the policy template **Grant limited access to non-private shared photos** with a friend group and individual photo.

```
permit (
 principal in PhotoFlash::FriendGroup::"Jane::MySchoolFriends",
 action in PhotoFlash::Action::"SharePhotoLimitedAccess",
 resource in PhotoFlash::Photo::"VacationPhoto94.jpg"
 );
```

The following example shows how you might create a template-linked policy that uses the policy template **Grant limited access to non-private shared photos** with a friend group and album.

```
permit (
 principal in PhotoFlash::FriendGroup::"Jane::MySchoolFriends",
 action in PhotoFlash::Action::"SharePhotoLimitedAccess",
 resource in PhotoFlash::Album::"Italy2023"
 );
```

The following example shows how you might create a template-linked policy that uses the policy template **Grant full access to non-private shared photos** with a friend group and an individual photo.

```
permit (
 principal in PhotoFlash::UserGroup::"Jane::MySchoolFriends",
 action in PhotoFlash::Action::"SharePhotoFullAccess",
 resource in PhotoFlash::Photo::"VacationPhoto94.jpg"
 );
```

The following example shows how you might create a template-linked policy that uses the policy template **Block user from an account**.

```
forbid(
 principal == PhotoFlash::User::"Bob",
 action,
 resource in PhotoFlash::Account::"Alice-account"
 );
```

DigitalPetStore examples

The DigitalPetStore sample policy store does not include any policy templates. You can view the policies included with the policy store by choosing **Policies** in the navigation pane on the left after creating the **DigitalPetStore** sample policy store.

TinyToDo examples

The following example shows how you might create a template-linked policy that uses the policy template that gives viewer access for an individual user and task list.

```
permit (
    principal == TinyTodo::User::"https://cognito-idp.us-east-1.amazonaws.com/us-
east-1_h2aKCU1ts|5ae0c4b1-6de8-4dff-b52e-158188686f31|bob",
    action in [TinyTodo::Action::"ReadList", TinyTodo::Action::"ListTasks"],
    resource == TinyTodo::List::"1"
);
```

The following example shows how you might create a template-linked policy that uses the policy template that gives editor access for an individual user and task list.

```
permit (
    principal == TinyTodo::User::"https://cognito-idp.us-east-1.amazonaws.com/us-
east-1_h2aKCU1ts|5ae0c4b1-6de8-4dff-b52e-158188686f31|bob",
    action in [
        TinyTodo::Action::"ReadList",
        TinyTodo::Action::"UpdateList",
        TinyTodo::Action::"ListTasks",
        TinyTodo::Action::"CreateTask",
        TinyTodo::Action::"UpdateTask",
        TinyTodo::Action::"DeleteTask"
    ],
    resource == TinyTodo::List::"1"
);
```

Secure your applications with identity sources and tokens

Secure you applications quickly by creating an *identity source* to represent an external identity provider (IdP) in Amazon Verified Permissions. Identity sources provide information from a user who authenticated with an IdP that has a trust relationship with your policy store. When your application makes an authorization request with a token from an identity source, your policy store can make authorization decisions from user properties and access permissions. You can add an Amazon Cognito user pool or a custom OpenID Connect (OIDC) IdP as your identity source.

You can use OpenID Connect (OIDC) identity providers (IdPs) with Verified Permissions. Your application can generate authorization requests with JSON web tokens (JWTs) generated by an OIDC-compliant identity provider. The user identity in the token is mapped to the principal ID. With ID tokens, Verified Permissions maps attribute claims to principal attributes. With Access tokens, these claims are mapped to context. With both token types, you can map a claim like `groups` to a principal group, and build policies that evaluate role-based access control (RBAC).

> **Note**
>
> Verified Permissions makes authorization decisions based on information from an IdP token but doesn't interact directly with the IdP in any way.

For a step-by-step walkthrough that builds authorization logic for Amazon API Gateway REST APIs using an Amazon Cognito user pool or OIDC identity provider, see Authorize API Gateway APIs using Amazon Verified Permissions with Amazon Cognito or bring your own identity provider on the *AWS Security Blog*.

Topics

- Choosing the right identity provider
- Working with Amazon Cognito identity sources
- Working with OIDC identity sources

Choosing the right identity provider

While Verified Permissions works with a variety of IdPs, consider the following when deciding which one to use in your application:

Use Amazon Cognito when:

- You're building new applications without existing identity infrastructure
- You want AWS-managed user pools with built-in security features
- You need social identity provider integration
- You want simplified token management

Use OIDC providers when:

- You have existing identity infrastructure (Auth0, Okta, Azure AD)
- You need to maintain centralized user management
- You have compliance requirements for specific IdPs

Working with Amazon Cognito identity sources

Verified Permissions works closely with Amazon Cognito user pools. Amazon Cognito JWTs have a predictable structure. Verified Permissions recognizes this structure and draws maximum benefit from the information that it contains. For example, you can implement a role-based access control (RBAC) authorization model with either ID tokens or access tokens.

A new Amazon Cognito user pools identity source requires the following information:

- The AWS Region.
- The user pool ID.
- The principal entity type that you want to associate with your identity source, for example `MyCorp::User`.
- The principal group entity type that you want to associate with your identity source, for example `MyCorp::UserGroup`.
- The client IDs from your user pool that you want to authorize to make requests to your policy store.

Because Verified Permissions only works with Amazon Cognito user pools in the same AWS account, you can't specify an identity source in another account. Verified Permissions sets the *entity prefix*—the identity-source identifier that you must reference in policies that act on user pool principals—to the ID of your user pool, for example `us-west-2_EXAMPLE`. In this case, you would reference a user in that user pool with ID `a1b2c3d4-5678-90ab-cdef-EXAMPLE22222` as `us-west-2_EXAMPLE|a1b2c3d4-5678-90ab-cdef-EXAMPLE22222`

User pool token *claims* can contain attributes, scopes, groups, client IDs, and custom data. Amazon Cognito JWTs have the ability to include a variety of information that can contribute to authorization decisions in Verified Permissions. These include:

1. Username and group claims with a `cognito:` prefix
2. Custom user attributes with a `custom: prefix`
3. Custom claims added at runtime
4. OIDC standard claims like `sub` and `email`

We cover these claims in detail, and how to manage them in Verified Permissions policies, in Mapping Amazon Cognito tokens to schema.

> **Important**
>
> Although you can revoke Amazon Cognito tokens before they expire, JWTs are considered to be stateless resources that are self-contained with a signature and validity. Services that conform with the JSON Web Token RFC 7519 are expected to validate tokens remotely and aren't required to validate them with the issuer. This means that it is possible for Verified Permissions to grant access based on a token that was revoked or issued for a user that was later deleted. To mitigate this risk, we recommend that you create your tokens with the shortest possible validity duration and revoke refresh tokens when you want to remove authorization to continue a user's session. For more information, see Ending user sessions with token revocation

This following example shows how you might create a policy that references some of the Amazon Cognito user pools claims associated with a principal.

```
permit(
     principal,
```

```
     action,
     resource == ExampleCo::Photo::"VacationPhoto94.jpg"
)
when {
     principal["cognito:username"]) == "alice" &&
     principal["custom:department"]) == "Finance"
};
```

This following example shows how you might create a policy that references a principal that's a user in a Cognito user pool. Note that the principal ID takes the form of "<userpool-id>|<sub>".

```
permit(
     principal == ExampleCo::User::"us-east-1_example|a1b2c3d4-5678-90ab-cdef-
EXAMPLE11111",
     action,
     resource == ExampleCo::Photo::"VacationPhoto94.jpg"
);
```

Cedar policies for user pool identity sources in Verified Permissions use a special syntax for claim names that contain characters other than alphanumeric and underscore (_). This includes user pool prefix claims that contain a `:` character, like `cognito:username` and `custom:department`. To write a policy condition that references the `cognito:username` or `custom:department` claim, write them as `principal["cognito:username"]` and `principal["custom:department"]`, respectively.

> **Note**
>
> If a token contains a claim with a `cognito:` or `custom:` prefix and a claim name with the literal value `cognito` or `custom`, an authorization request with IsAuthorizedWithToken will fail with a `ValidationException`.

For more information about mapping claims, see Mapping Amazon Cognito tokens to schema. For more information about authorization for Amazon Cognito users, see Authorization with Amazon Verified Permissions in the *Amazon Cognito Developer Guide*.

Topics

- Creating Amazon Verified Permissions Amazon Cognito identity sources

- Editing Amazon Verified Permissions Amazon Cognito identity sources
- Mapping Amazon Cognito tokens to schema
- Client and audience validation for Amazon Cognito

Creating Amazon Verified Permissions Amazon Cognito identity sources

The following procedure adds an identity source to an existing policy store.

You can also create an identity source when you create a new policy store in the Verified Permissions console. In this process, you can automatically import the claims in your identity source tokens into entity attributes. Choose the **Guided setup** or **Set up with API Gateway and an identity provider** option. These options also create initial policies.

> **Note**
>
> **Identity sources** is not available in the navigation pane on the left until you have created a policy store. Identity sources that you create are associated with the current policy store.

You can leave out the principal entity type when you create an identity source with create-identity-source in the AWS CLI or CreateIdentitySource in the Verified Permissions API. However, a blank entity type creates an identity source with an entity type of `AWS::Cognito`. This entity name isn't compatible with policy store schema. To integrate Amazon Cognito identities with your policy store schema, you must set the principal entity type to a supported policy store entity.

AWS Management Console

To create an Amazon Cognito user pools identity source

1. Open the Verified Permissions console. Choose your policy store.
2. In the navigation pane on the left, choose **Identity sources**.
3. Choose **Create identity source**.
4. In **Cognito user pool details**, select the AWS Region and enter the **User pool ID** for your identity source.
5. In **Principal configuration**, for **Principal type**, choose the entity type for principals from this source. Identities from the connected Amazon Cognito user pools will be mapped to the selected principal type.

6. In **Group configuration**, select **Use Cognito group** if you want to map the user pool `cognito:groups` claim. Choose an entity type that is a parent of the principal type.
7. In **Client application validation**, choose whether to validate client application IDs.
   - To validate client application IDs, choose **Only accept tokens with matching client application IDs**. Choose **Add new client application ID** for each client application ID to validate. To remove a client application ID that has been added, choose **Remove** next to the client application ID.
   - Choose **Do not validate client application IDs** if you do not want to validate client application IDs.
8. Choose **Create identity source**.
9. (Optional) If your policy store has a schema, before you can reference attributes you extract from identity or access tokens in your Cedar policies, you must update your schema to make Cedar aware of the type of principal that your identity source creates. That addition to the schema must include the attributes that you want to reference in your Cedar policies. For more information about mapping Amazon Cognito token attributes to Cedar principal attributes, see Mapping Amazon Cognito tokens to schema.

   > **Note**
   >
   > When you create an API-linked policy store or use **Set up with API Gateway and an identity provider** when creating policy stores, Verified Permissions queries your user pool for user attributes and creates a schema where your principal type is populated with user pool attributes.

10. Create policies that use information from the tokens to make authorization decisions. For more information, see Creating Amazon Verified Permissions static policies.

Now that you've created an identity source, updated the schema, and created policies, use `IsAuthorizedWithToken` to have Verified Permissions make authorization decisions. For more information, see IsAuthorizedWithToken in the *Amazon Verified Permissions API reference guide*.

AWS CLI

To create an Amazon Cognito user pools identity source

You can an create an identity source by using the CreateIdentitySource operation. The following example creates an identity source that can access authenticated identities from a Amazon Cognito user pool.

1. Create a `config.txt` file that contains the following details of the Amazon Cognito user pool for use by the `--configuration` parameter in the `create-identity-source` command.

```
{
    "cognitoUserPoolConfiguration": {
        "userPoolArn": "arn:aws:cognito-idp:us-west-2:123456789012:userpool/us-
west-2_1a2b3c4d5",
        "clientIds":["a1b2c3d4e5f6g7h8i9j0kalbmc"],
        "groupConfiguration": {
              "groupEntityType": "MyCorp::UserGroup"
        }
    }
}
```

2. Run the following command to create an Amazon Cognito identity source.

```
$ aws verifiedpermissions create-identity-source \
    --configuration file://config.txt \
    --principal-entity-type "User" \
    --policy-store-id 123456789012
{
    "createdDate": "2023-05-19T20:30:28.214829+00:00",
    "identitySourceId": "ISEXAMPLEabcdefg111111",
    "lastUpdatedDate": "2023-05-19T20:30:28.214829+00:00",
    "policyStoreId": "PSEXAMPLEabcdefg111111"
}
```

3. (Optional) If your policy store has a schema, before you can reference attributes you extract from identity or access tokens in your Cedar policies, you must update your schema to make Cedar aware of the type of principal that your identity source creates. That addition to the schema must include the attributes that you want to reference in your Cedar policies. For more information about mapping Amazon Cognito token attributes to Cedar principal attributes, see Mapping Amazon Cognito tokens to schema.

> **Note**
>
> When you create an API-linked policy store or use **Set up with API Gateway and an identity provider** when creating policy stores, Verified Permissions queries your user pool for user attributes and creates a schema where your principal type is populated with user pool attributes.

4. Create policies that use information from the tokens to make authorization decisions. For more information, see Creating Amazon Verified Permissions static policies.

Now that you've created an identity source, updated the schema, and created policies, use `IsAuthorizedWithToken` to have Verified Permissions make authorization decisions. For more information, see IsAuthorizedWithToken in the *Amazon Verified Permissions API reference guide*.

For more information about using Amazon Cognito access and identity tokens for authenticated users in Verified Permissions, see Authorization with Amazon Verified Permissions in the *Amazon Cognito Developer Guide*.

Editing Amazon Verified Permissions Amazon Cognito identity sources

You can edit some parameters of your identity source after you create it. You can't change the type of identity source, you have to delete the identity source and create a new one to switch from Amazon Cognito to OIDC or OIDC to Amazon Cognito. If your policy store schema matches your identity source attributes, note that you must update your schema separately to reflect the changes that you make to your identity source.

AWS Management Console

To update an Amazon Cognito identity source

1. Open the Verified Permissions console. Choose your policy store.
2. In the navigation pane on the left, choose **Identity sources**.
3. Choose the ID of the identity source to edit.
4. Choose **Edit**.

5. In **Cognito user pool details**, select the AWS Region and type the **User pool ID** for your identity source.
6. In **Principal details**, you can update the **Principal type** for the identity source. Identities from the connected Amazon Cognito user pools will be mapped to the selected principal type.
7. In **Group configuration**, select **Use Cognito groups** if you want to map the user pool `cognito:groups` claim. Choose an entity type that is a parent of the principal type.
8. In **Client application validation**, choose whether to validate client application IDs.
   - To validate client application IDs, choose **Only accept tokens with matching client application IDs**. Choose **Add new client application ID** for each client application ID to validate. To remove a client application ID that has been added, choose **Remove** next to the client application ID.
   - Choose **Do not validate client application IDs** if you do not want to validate client application IDs.
9. Choose **Save changes**.
10. If you changed the principal type for the identity source, you must update your schema to correctly reflect the updated principal type.

You can delete an identity source by choosing the radio button next to an identity source and then choosing **Delete identity source**. Type `delete` in the text box and then choose **Delete identity source** to confirm deleting the identity source.

AWS CLI

To update an Amazon Cognito identity source

You can update an identity source by using the UpdateIdentitySource operation. The following example updates the specified identity source to use a different Amazon Cognito user pool.

1. Create a `config.txt` file that contains the following details of the Amazon Cognito user pool for use by the `--configuration` parameter in the `update-identity-source` command.

```
{
    "cognitoUserPoolConfiguration": {
        "userPoolArn": "arn:aws:cognito-idp:us-west-2:123456789012:userpool/us-
west-2_1a2b3c4d5",
```

```
        "clientIds":["a1b2c3d4e5f6g7h8i9j0kalbmc"],
        "groupConfiguration": {
              "groupEntityType": "MyCorp::UserGroup"
        }
    }
}
```

2. Run the following command to update an Amazon Cognito identity source.

```
$ aws verifiedpermissions update-identity-source \
    --update-configuration file://config.txt \
    --policy-store-id 123456789012
{
    "createdDate": "2023-05-19T20:30:28.214829+00:00",
    "identitySourceId": "ISEXAMPLEabcdefg111111",
    "lastUpdatedDate": "2023-05-19T20:30:28.214829+00:00",
    "policyStoreId": "PSEXAMPLEabcdefg111111"
}
```

> **Note**
>
> If you change the principal type for the identity source, you must update your schema to correctly reflect the updated principal type.

Mapping Amazon Cognito tokens to schema

You might find that you want to add an identity source to a policy store and map provider claims, or tokens, to your policy store schema. You can automate this process, by using the Guided setup to create your policy store with an identity source, or update your schema manually after the policy store is created. Once you have mapped the tokens to the schema you can create policies that reference them.

This section of the user guide has the following information:

- When you can automatically populate attributes to a policy store schema
- How to use Amazon Cognito token claims in your Verified Permissions policies
- How to manually build a schema for an identity source

API-linked policy stores and policy stores with an identity source that were created through Guided setup don't require manual mapping of identity (ID) token attributes to schema. You can provide Verified Permissions with the attributes in your user pool and create a schema that is populated with user attributes. In ID token authorization, Verified Permissions maps claims to attributes of a principal entity. You might need to manually map Amazon Cognito tokens to your schema in the following conditions:

- You created an empty policy store or policy store from a sample.
- You want to extend your use of access tokens beyond role-based access control (RBAC).
- You create policy stores with the Verified Permissions REST API, an AWS SDK, or the AWS CDK.

To use Amazon Cognito as an identity source in your Verified Permissions policy store, you must have provider attributes in your schema. The schema is fixed and must correspond to the entities that provider tokens create in IsAuthorizedWithToken or BatchIsAuthorizedWithToken API requests. If you created your policy store in a way that automatically populates your schema from provider information in an ID token, you're ready to write policies. If you create a policy store without a schema for your identity source, you must add provider attributes to the schema that match the entities created using API requests. Then you can write policies using attributes from the provider token.

For more information about using Amazon Cognito ID and access tokens for authenticated users in Verified Permissions, see Authorization with Amazon Verified Permissions in the *Amazon Cognito Developer Guide*.

Topics

Mapping ID tokens to schema

Verified Permissions processes ID token claims as the attributes of the user: their names and titles, their group membership, their contact information. ID tokens are most useful in an *attribute-based access control* (ABAC) authorization model. When you want Verified Permissions to analyze access to resources based on who's making the request, choose ID tokens for your identity source.

Amazon Cognito ID tokens work with most OIDC relying-party libraries. They extend the features of OIDC with additional claims. Your application can authenticate the user with Amazon Cognito user pools authentication API operations, or with the user pool hosted UI. For more information, see Using the API and endpoints in the *Amazon Cognito Developer Guide*.

Useful claims in Amazon Cognito ID tokens

`cognito:username` and `preferred_username`

Variants of the user's username.

`sub`

The user's unique user identifier (UUID)

Claims with a `custom:` prefix

A prefix for custom user pool attributes like `custom:employmentStoreCode`.

Standard claims

Standard OIDC claims like `email` and `phone_number`. For more information, see Standard claims in *OpenID Connect Core 1.0 incorporating errata set 2*.

`cognito:groups`

A user's group memberships. In an authorization model based on role-based access control (RBAC), this claim presents the roles that you can evaluate in your policies.

Transient claims

Claims that aren't a property of the user, but are added at runtime by a user pool Pre token generation Lambda trigger. Transient claims resemble standard claims but are outside the standard, for example `tenant` or `department`.

In policies that reference Amazon Cognito attributes that have a `:` separator, reference the attributes in the format `principal["cognito:username"]`. The roles claim `cognito:groups` is an exception to this rule. Verified Permissions maps the contents of this claim to parent entities of the user entity.

For more information about the structure of ID tokens from Amazon Cognito user pools, see Using the ID token in the *Amazon Cognito Developer Guide*.

The following example ID token has each of the four types of attributes. It includes the Amazon Cognito-specific claim `cognito:username`, the custom claim `custom:employmentStoreCode`, the standard claim `email`, and the transient claim `tenant`.

```
{
    "sub": "91eb4550-XXX",
    "cognito:groups": [
        "Store-Owner-Role",
        "Customer"
    ],
    "email_verified": true,
    "clearance": "confidential",
    "iss": "https://cognito-idp.us-east-2.amazonaws.com/us-east-2_EXAMPLE",
    "cognito:username": "alice",
    "custom:employmentStoreCode": "petstore-dallas",
    "origin_jti": "5b9f50a3-05da-454a-8b99-b79c2349de77",
    "aud": "1example23456789",
    "event_id": "0ed5ad5c-7182-4ecf-XXX",
    "token_use": "id",
    "auth_time": 1687885407,
    "department": "engineering",
    "exp": 1687889006,
    "iat": 1687885407,
    "tenant": "x11app-tenant-1",
    "jti": "a1b2c3d4-e5f6-a1b2-c3d4-TOKEN1111111",
    "email": "alice@example.com"
}
```

When you create an identity source with your Amazon Cognito user pool, you specify the type of principal entity that Verified Permissions generates in authorization requests with `IsAuthorizedWithToken`. Your policies can then test attributes of that principal as part of evaluating that request. Your schema defines the principal type and attributes for an identity source, and then you can reference them in your Cedar policies.

You also specify the type of group entity that you want to derive from the ID token groups claim. In authorization requests, Verified Permissions maps each member of the groups claim to that group entity type. In policies, you can reference that group entity as the principal.

The following example shows how to reflect the attributes from the example identity token in your Verified Permissions schema. For more information about editing your schema, see Editing policy

store schemas. If your identity source configuration specifies the principal type `User`, then you can include something similar to the following example to make those attributes available to Cedar.

```
"User": {
   "shape": {
      "type": "Record",
      "attributes": {
         "cognito:username": {
            "type": "String",
            "required": false
         },
         "custom:employmentStoreCode": {
            "type": "String",
            "required": false
         },
         "email": {
            "type": "String"
         },
         "tenant": {
            "type": "String",
            "required": true
         }
      }
   }
}
```

For an example policy that will validate against this schema, see Reflects Amazon Cognito ID token attributes.

Mapping access tokens

Verified Permissions processes access-token claims other than the groups claim as attributes of the action, or *context attributes*. Along with group membership, the access tokens from your IdP might contain information about API access. Access tokens are useful in authorization models that use role-based access control (RBAC). Authorization models that rely on access-token claims other than group membership require additional effort in schema configuration.

Amazon Cognito access tokens have claims that can be used for authorization:

Useful claims in Amazon Cognito access tokens

`client_id`

The ID of the client application of an OIDC relying party. With the client ID, Verified Permissions can verify that the authorization request comes from a permitted client for the policy store. In machine-to-machine (M2M) authorization, the requesting system authorizes a request with a client secret and provides the client ID and scopes as evidence of authorization.

`scope`

The OAuth 2.0 scopes that represent the access permissions of the bearer of the token.

`cognito:groups`

A user's group memberships. In an authorization model based on role-based access control (RBAC), this claim presents the roles that you can evaluate in your policies.

Transient claims

Claims that aren't an access permission, but are added at runtime by a user pool Pre token generation Lambda trigger. Transient claims resemble standard claims but are outside the standard, for example `tenant` or `department`. Customization of access tokens adds cost to your AWS bill.

For more information about the structure of access tokens from Amazon Cognito user pools, see Using the access token in the *Amazon Cognito Developer Guide*.

An Amazon Cognito access token is mapped to a context object when passed to Verified Permissions. Attributes of the access token can be referenced using `context.token.`*`attribute_name`*. The following example access token includes both the `client_id` and `scope` claims.

```
{
    "sub": "91eb4550-9091-708c-a7a6-9758ef8b6b1e",
    "cognito:groups": [
        "Store-Owner-Role",
        "Customer"
    ],
    "iss": "https://cognito-idp.us-east-2.amazonaws.com/us-east-2_EXAMPLE",
    "client_id": "1example23456789",
    "origin_jti": "a1b2c3d4-e5f6-a1b2-c3d4-TOKEN1111111",
    "event_id": "bda909cb-3e29-4bb8-83e3-ce6808f49011",
```

```
    "token_use": "access",
    "scope": "MyAPI/mydata.write",
    "auth_time": 1688092966,
    "exp": 1688096566,
    "iat": 1688092966,
    "jti": "a1b2c3d4-e5f6-a1b2-c3d4-TOKEN2222222",
    "username": "alice"
}
```

The following example shows how to reflect the attributes from the example access token in your Verified Permissions schema. For more information about editing your schema, see Editing policy store schemas.

```
{
   "MyApplication": {
      "actions": {
         "Read": {
            "appliesTo": {
               "context": {
                  "type": "ReusedContext"
               },
               "resourceTypes": [
                  "Application"
               ],
               "principalTypes": [
                  "User"
               ]
            }
         }
      },
      ...
      ...
      "commonTypes": {
         "ReusedContext": {
            "attributes": {
               "token": {
                  "type": "Record",
                  "attributes": {
                     "scope": {
                        "type": "Set",
                        "element": {
                           "type": "String"
                        }
```

```
                    },
                    "client_id": {
                       "type": "String"
                    }
                 }
              }
           },
           "type": "Record"
        }
     }
  }
}
```

For an example policy that will validate against this schema, see Reflects Amazon Cognito access token attributes.

Alternative notation for Amazon Cognito colon-delimited claims

At the time that Verified Permissions launched, the recommended schema for Amazon Cognito token claims like `cognito:groups` and `custom:store` converted these colon-delimited strings to use the `.` character as a hierarchy delimiter. This format is called *dot notation*. For example, a reference to `cognito:groups` became `principal.cognito.groups` in your policies. Although you can continue to use this format, we recommend that you build your schema and policies with bracket notation. In this format, a reference to `cognito:groups` becomes `principal["cognito:groups"]` in your policies. Automatically-generated schemas for user pool ID tokens from the Verified Permissions console use bracket notation.

You can continue to use dot notation in manually-built schema and policies for Amazon Cognito identity sources. You can't use dot notation with `:` or any other non-alphanumeric characters in schema or policies for any other type of OIDC IdP.

A schema for dot notation nests each instance of a `:` character as a child of the `cognito` or `custom` initial phrase, as shown in the following example:

```
"CognitoUser": {
   "shape": {
      "type": "Record",
      "attributes": {
         "cognito": {
            "type": "Record",
            "required": true,
```

```
            "attributes": {
               "username": {
                  "type": "String",
                  "required": true
               }
            }
         },
         "custom": {
            "type": "Record",
            "required": true,
            "attributes": {
               "employmentStoreCode": {
                  "type": "String",
                  "required": true
               }
            }
         },
         "email": {
            "type": "String"
         },
         "tenant": {
            "type": "String",
            "required": true
         }
      }
   }
}
```

For an example policy that will validate against this schema and use dot notation, see Uses dot notation to reference attributes.

Things to know about schema mapping

Attribute mapping differs between token types

In access token authorization, Verified Permissions maps claims to context. In ID token authorization, Verified Permissions maps claims to principal attributes. For policy stores that you create in the Verified Permissions console, only **empty** and **sample** policy stores leave you with no identity source and require you to populate your schema with user pool attributes for ID token authorization. Access token authorization is based on role-based access control (RBAC) with group-membership claims and doesn't automatically map other claims to the policy store schema.

Identity source attributes aren't required

When you create an identity source in the Verified Permissions console, no attributes are marked as required. This prevents missing claims from causing validation errors in authorization requests. You can set attributes to required as needed, but they must be present in all authorization requests.

RBAC doesn't require attributes in schema

Schemas for identity sources depend on the entity associations that you make when you add your identity source. An identity source maps one claim to a user entity type, and one claim to a group entity type. These entity mappings are the core of an identity-source configuration. With this minimum information, you can write policies that perform authorization actions for specific users and specific groups that users might be members of, in a role-based access control (RBAC) model. The addition of token claims to the schema extends the authorization scope of your policy store. User attributes from ID tokens have information about users that can contribute to attribute-based access control (ABAC) authorization. Context attributes from access tokens have information like OAuth 2.0 scopes that can contribute additional access-control information from your provider, but require additional schema modifications.

The **Set up with API Gateway and an identity provider** and **Guided setup** options in the Verified Permissions console assign ID token claims to the schema. This isn't the case for access token claims. To add non-group access-token claims to your schema, you must edit your schema in JSON mode and add commonTypes attributes. For more information, see Mapping access tokens.

Choose a token type

The way that your policy store works with your identity source depends on a key decision in identity-source configuration: whether you will process ID or access tokens. With an Amazon Cognito identity provider, you have the choice of token type when you create an API-linked policy store. When you create an API-linked policy store, you must choose whether you want to set up authorization for ID or access tokens. This information affects the schema attributes that Verified Permissions applies to your policy store, and the syntax of the Lambda authorizer for your API Gateway API. Especially if you wish to benefit from the automatic mapping of ID token claims to attributes in the Verified Permissions console, decide early about the token type that you want to process before you create your identity source. Changing the token type requires significant effort to refactor your policies and schema. The following topics describe the use of ID and access tokens with policy stores.

Cedar parser requires brackets for some characters

Policies typically reference schema attributes in a format like `principal.username`. In the case of most non-alphanumeric characters like `:`, `.`, or `/` that might appear in token claim names,

Verified Permissions can't parse a condition value like `principal.cognito:username` or `context.ip-address`. You must instead format these conditions with bracket notation in the format `principal["cognito:username"]` or `context["ip-address"]`, respectively. The underscore character _ is a valid character in claim names, and the only non-alphanumeric exception to this requirement.

A partial example schema for a principal attribute of this type looks like the following:

```
"User": {
   "shape": {
      "type": "Record",
      "attributes": {
         "cognito:username": {
            "type": "String",
            "required": true
         },
         "custom:employmentStoreCode": {
            "type": "String",
            "required": true,
         },
         "email": {
            "type": "String",
            "required": false
         }
      }
   }
}
```

A partial example schema for a context attribute of this type looks like the following:

```
"GetOrder": {
   "memberOf": [],
   "appliesTo": {
      "resourceTypes": [
         "Order"
      ],
      "context": {
         "type": "Record",
         "attributes": {
            "ip-address": {
               "required": false,
               "type": "String"
```

```
            }
    }
    },
       "principalTypes": [
          "User"
       ]
    }
}
```

For an example policy that will validate against this schema, see Uses bracket notation to reference token attributes.

Client and audience validation for Amazon Cognito

When you add an identity source to a policy store, Verified Permissions has configuration options that verify that ID and access tokens are being used as intended. This validation happens in the processing of `IsAuthorizedWithToken` and `BatchIsAuthorizedWithToken` API requests. The behavior differs between ID and access tokens, and between Amazon Cognito and OIDC identity sources. With Amazon Cognito user pools providers, Verified Permissions can validate the client ID in both ID and access tokens. With OIDC providers, Verified Permissions can validate the client ID in ID tokens, and the audience in access tokens.

A *client ID* is an identifier associated with the identity provider instance that your application uses, for example `1example23456789`. An *audience* is a URL path associated with the intended *relying party*, or destination, of the access token, for example `https://mytoken.example.com`. When using access tokens, the `aud` claim is always associated with the audience.

Amazon Cognito ID tokens have an `aud` claim that contains the app client ID. Access tokens have a `client_id` claim that also contains the app client ID.

When you enter one or more values for **Client application validation** in your identity source, Verified Permissions compares this list of app client IDs to the ID token `aud` claim or the access token `client_id` claim. Verified Permissions doesn't validate a relying-party audience URL for Amazon Cognito identity sources.

Client-side authorization for JWTs

You might want to process JSON web tokens in your application and pass their claims to Verified Permissions without using a policy store identity source. You can extract your entity attributes from a JSON Web Token (JWT) and parse it into Verified Permissions.

This example shows how you might call Verified Permissions from an application using a JWT.[1]

```
async function authorizeUsingJwtToken(jwtToken) {

    const payload = await verifier.verify(jwtToken);

    let principalEntity = {
        entityType: "PhotoFlash::User", // the application needs to fill in the
 relevant user type
        entityId: payload["sub"], // the application need to use the claim that
 represents the user-id
    };
    let resourceEntity = {
        entityType: "PhotoFlash::Photo", //the application needs to fill in the
 relevant resource type
        entityId: "jane_photo_123.jpg", // the application needs to fill in the
 relevant resource id
    };
    let action = {
        actionType: "PhotoFlash::Action", //the application needs to fill in the
 relevant action id
        actionId: "GetPhoto", //the application needs to fill in the relevant action
 type
    };
    let entities = {
        entityList: [],
    };
    entities.entityList.push(...getUserEntitiesFromToken(payload));
    let policyStoreId = "PSEXAMPLEabcdefg111111"; // set your own policy store id

    const authResult = await client
        .isAuthorized({
        policyStoreId: policyStoreId,
        principal: principalEntity,
        resource: resourceEntity,
        action: action,
        entities,
        })
        .promise();

    return authResult;

}
```

```
function getUserEntitiesFromToken(payload) {
  let attributes = {};
  let claimsNotPassedInEntities = ['aud', 'sub', 'exp', 'jti', 'iss'];
  Object.entries(payload).forEach(([key, value]) => {
    if (claimsNotPassedInEntities.includes(key)) {
        return;
    }
    if (Array.isArray(value)) {
      var attibuteItem = [];
      value.forEach((item) => {
        attibuteItem.push({
          string: item,
        });
      });
      attributes[key] = {
        set: attibuteItem,
      };
    } else if (typeof value === 'string') {
      attributes[key] = {
        string: value,
      }
    } else if (typeof value === 'bigint' || typeof value ==='number') {
        attributes[key] = {
            long: value,
          }
    } else if (typeof value === 'boolean') {
        attributes[key] = {
            boolean: value,
       }
    }

  });

  let entityItem = {
    attributes: attributes,
    identifier: {
      entityType: "PhotoFlash::User",
      entityId: payload["sub"], // the application needs to use the claim that
 represents the user-id
    }
  };
  return [entityItem];
}
```

[1] This code example uses the aws-jwt-verify library for verifying JWTs signed by OIDC-compatible IdPs.

Working with OIDC identity sources

You can also configure any compliant OpenID Connect (OIDC) IdP as the identity source of a policy store. OIDC providers are similar to Amazon Cognito user pools: they produce JWTs as the product of authentication. To add an OIDC provider, you must provide an issuer URL

A new OIDC identity source requires the following information:

- The issuer URL. Verified Permissions must be able to discover a `.well-known/openid-configuration` endpoint at this URL.
- CNAME records that don't include wild cards. For example, a `.example.com` can't be mapped to `*.example.net`. Conversely, `*.example.com` can't be mapped to `a.example.net`.
- The token type that you want to use in authorization requests. In this case, you chose **Identity token**.
- The user entity type that you want to associate with your identity source, for example `MyCorp::User`.
- The group entity type that you want to associate with your identity source, for example `MyCorp::UserGroup`.
- An example ID token, or a definition of the claims in the ID token.
- The prefix that you want to apply to user and group entity IDs. In the CLI and API, you can choose this prefix. In policy stores that you create with the **Set up with API Gateway and an identity provider** or **Guided setup** option, Verified Permissions assigns a prefix of the issuer name minus `https://`, for example `MyCorp::User::"auth.example.com|a1b2c3d4-5678-90ab-cdef-EXAMPLE11111"`.

For more information about using API operations to authorize requests from OIDC sources, see Available API operations for authorization.

This following example shows how you might create a policy that permits access to year-end reports for employees in the accounting department, have a confidential classification, and aren't in a satellite office. Verified Permissions derives these attributes from the claims in the principal's ID token.

Note that when referencing a group in the principal, you must use the `in` operator for the policy to be evaluated correctly.

```
permit(
     principal in MyCorp::UserGroup::"MyOIDCProvider|Accounting",
     action,
     resource in MyCorp::Folder::"YearEnd2024"
) when {
     principal.jobClassification == "Confidential" &&
     !(principal.location like "SatelliteOffice*")
};
```

Topics

- Creating Amazon Verified Permissions OIDC identity sources
- Editing Amazon Verified Permissions OIDC identity sources
- Mapping OIDC tokens to schema
- Client and audience validation for OIDC providers

Creating Amazon Verified Permissions OIDC identity sources

The following procedure adds an identity source to an existing policy store.

You can also create an identity source when you create a new policy store in the Verified Permissions console. In this process, you can automatically import the claims in your identity source tokens into entity attributes. Choose the **Guided setup** or **Set up with API Gateway and an identity provider** option. These options also create initial policies.

> **Note**
>
> **Identity sources** is not available in the navigation pane on the left until you have created a policy store. Identity sources that you create are associated with the current policy store.

You can leave out the principal entity type when you create an identity source with create-identity-source in the AWS CLI or CreateIdentitySource in the Verified Permissions API. However, a blank entity type creates an identity source with an entity type of `AWS::Cognito`. This entity name isn't compatible with policy store schema. To integrate Amazon Cognito identities with your policy store schema, you must set the principal entity type to a supported policy store entity.

AWS Management Console

To create an OpenID Connect (OIDC) identity source

1. Open the Verified Permissions console. Choose your policy store.
2. In the navigation pane on the left, choose **Identity sources**.
3. Choose **Create identity source**.
4. Choose **External OIDC provider**.
5. In **Issuer URL**, enter the URL of your OIDC issuer. This is the service endpoint that provides the authorization server, signing keys, and other information about your provider, for example `https://auth.example.com`. Your issuer URL must host an OIDC discovery document at `/.well-known/openid-configuration`.
6. In **Token type**, choose the type of OIDC JWT that you want your application to submit for authorization. For more information, see Mapping OIDC tokens to schema.
7. In **Map token claims to schema entities**, choose a **User entity** and **User claim** for the identity source. The **User entity** is an entity in your policy store that you want to refer to users from your OIDC provider. The **User claim** is a claim, typically `sub`, from your ID or access token that holds the unique identifier for the entity to be evaluated. Identities from the connected OIDC IdP will be mapped to the selected principal type.
8. (Optional) In **Map token claims to schema entities**, choose a **Group entity** and **Group claim** for the identity source. The **Group entity** is a parent of the **User entity**. Group claims get mapped to this entity. The **Group claim** is a claim, typically `groups`, from your ID or access token that contains a string, JSON, or space-delimited string of user-group names for the entity to be evaluated. Identities from the connected OIDC IdP will be mapped to the selected principal type.
9. In **validation - optional**, enter the client IDs or audience URLs that you want your policy store to accept in authorization requests, if any.
10. Choose **Create identity source**.
11. (Optional) If your policy store has a schema, before you can reference attributes that you extract from identity or access tokens in your Cedar policies, you must update your schema to make Cedar aware of the type of principal that your identity source creates. That addition to the schema must include the attributes that you want to reference in your Cedar policies. For more information about mapping OIDC token attributes to Cedar principal attributes, see Mapping OIDC tokens to schema.

12. Create policies that use information from the tokens to make authorization decisions. For more information, see Creating Amazon Verified Permissions static policies.

Now that you've created an identity source, updated the schema, and created policies, use `IsAuthorizedWithToken` to have Verified Permissions make authorization decisions. For more information, see IsAuthorizedWithToken in the *Amazon Verified Permissions API reference guide*.

AWS CLI

To create an OIDC identity source

You can an create an identity source by using the CreateIdentitySource operation. The following example creates an identity source that can access authenticated identities from a an OIDC identity provider(IdP).

1. Create a `config.txt` file that contains the following details of an OIDC IdP for use by the `--configuration` parameter of the `create-identity-source` command.

```
{
    "openIdConnectConfiguration": {
        "issuer": "https://auth.example.com",
        "tokenSelection": {
                "identityTokenOnly": {
                        "clientIds":["1example23456789"],
                        "principalIdClaim": "sub"
                },
        },
        "entityIdPrefix": "MyOIDCProvider",
        "groupConfiguration": {
              "groupClaim": "groups",
              "groupEntityType": "MyCorp::UserGroup"
        }
    }
}
```

2. Run the following command to create an OIDC identity source.

```
$ aws verifiedpermissions create-identity-source \
    --configuration file://config.txt \
    --principal-entity-type "User" \
    --policy-store-id 123456789012
```

```
{
    "createdDate": "2023-05-19T20:30:28.214829+00:00",
    "identitySourceId": "ISEXAMPLEabcdefg111111",
    "lastUpdatedDate": "2023-05-19T20:30:28.214829+00:00",
    "policyStoreId": "PSEXAMPLEabcdefg111111"
}
```

3. (Optional) If your policy store has a schema, before you can reference attributes that you extract from identity or access tokens in your Cedar policies, you must update your schema to make Cedar aware of the type of principal that your identity source creates. That addition to the schema must include the attributes that you want to reference in your Cedar policies. For more information about mapping OIDC token attributes to Cedar principal attributes, see Mapping OIDC tokens to schema.
4. Create policies that use information from the tokens to make authorization decisions. For more information, see Creating Amazon Verified Permissions static policies.

Now that you've created an identity source, updated the schema, and created policies, use `IsAuthorizedWithToken` to have Verified Permissions make authorization decisions. For more information, see IsAuthorizedWithToken in the *Amazon Verified Permissions API reference guide*.

Editing Amazon Verified Permissions OIDC identity sources

You can edit some parameters of your identity source after you create it. You can't change the type of identity source, you have to delete the identity source and create a new one to switch from Amazon Cognito to OIDC or OIDC to Amazon Cognito. If your policy store schema matches your identity source attributes, note that you must update your schema separately to reflect the changes that you make to your identity source.

AWS Management Console

To update an OIDC identity source

1. Open the Verified Permissions console. Choose your policy store.
2. In the navigation pane on the left, choose **Identity sources**.
3. Choose the ID of the identity source to edit.
4. Choose **Edit**.

5. In **OIDC provider details**, change the **Issuer URL** as needed.
6. In **Map token claims to schema attributes**, change the associations between user and group claims and policy store entity types, as needed. After you change entity types, you must update your policies and schema attributes to apply to the new entity types.
7. In **Audience validation**, add or remove audience values that you want to enforce.
8. Choose **Save changes**.

You can delete an identity source by choosing the radio button next to an identity source and then choosing **Delete identity source**. Type `delete` in the text box and then choose **Delete identity source** to confirm deleting the identity source.

AWS CLI

To update an OIDC identity source

You can update an identity source by using the UpdateIdentitySource operation. The following example updates the specified identity source to use a different OIDC provider.

1. Create a `config.txt` file that contains the following details of an OIDC IdP for use by the `--configuration` parameter of the `update-identity-source` command.

```
{
    "openIdConnectConfiguration": {
        "issuer": "https://auth2.example.com",
        "tokenSelection": {
                "identityTokenOnly": {
                        "clientIds":["2example10111213"],
                        "principalIdClaim": "sub"
                },
        },
        "entityIdPrefix": "MyOIDCProvider",
        "groupConfiguration": {
              "groupClaim": "groups",
              "groupEntityType": "MyCorp::UserGroup"
        }
    }
}
```

2. Run the following command to update an OIDC identity source.

```
$ aws verifiedpermissions update-identity-source \
```

```
    --update-configuration file://config.txt \
    --policy-store-id 123456789012
{
    "createdDate": "2023-05-19T20:30:28.214829+00:00",
    "identitySourceId": "ISEXAMPLEabcdefg111111",
    "lastUpdatedDate": "2023-05-19T20:30:28.214829+00:00",
    "policyStoreId": "PSEXAMPLEabcdefg111111"
}
```

> **Note**
>
> If you change the principal type for the identity source, you must update your schema to correctly reflect the updated principal type.

Mapping OIDC tokens to schema

You might find that you want to add an identity source to a policy store and map provider claims, or tokens, to your policy store schema. You can automate this process, by using the Guided setup to create your policy store with an identity source, or update your schema manually after the policy store is created. Once you have mapped the tokens to the schema you can create policies that reference them.

This section of the user guide has the following information:

- When you can automatically populate attributes to a policy store schema
- How to manually build a schema for an identity source

API-linked policy stores and policy stores with an identity source that were created through Guided setup don't require manual mapping of identity (ID) token attributes to schema. You can provide Verified Permissions with the attributes in your user pool and create a schema that is populated with user attributes. In ID token authorization, Verified Permissions maps claims to attributes of a principal entity.

To use an OIDC identity provider (IdP) as an identity source in your Verified Permissions policy store, you must have provider attributes in your schema. The schema is fixed and must correspond to the entities that provider tokens create in IsAuthorizedWithToken or BatchIsAuthorizedWithToken API requests. If you created your policy store in a way that

automatically populates your schema from provider information in an ID token, you're ready to write policies. If you create a policy store without a schema for your identity source, you must add provider attributes to the schema that match the entities created using API requests. Then you can write policies using attributes from the provider token.

Topics

Mapping ID tokens to schema

Verified Permissions processes ID token claims as the attributes of the user: their names and titles, their group membership, their contact information. ID tokens are most useful in an *attribute-based access control* (ABAC) authorization model. When you want Verified Permissions to analyze access to resources based on who's making the request, choose ID tokens for your identity source.

Working with ID tokens from an OIDC provider is much the same as working with Amazon Cognito ID tokens. The difference is in the claims. Your IdP might present standard OIDC attributes, or have a custom schema. When you create a new policy store in the Verified Permissions console, you can add an OIDC identity source with an example ID token, or you can manually map token claims to user attributes. Because Verified Permissions isn't aware of the attribute schema of your IdP, you must provide this information.

For more information, see Creating Verified Permissions policy stores.

The following is an example schema for a policy store with an OIDC identity source.

```
"User": {
   "shape": {
      "type": "Record",
      "attributes": {
         "email": {
            "type": "String"
         },
         "email_verified": {
            "type": "Boolean"
         },
         "name": {
```

```
            "type": "String",
            "required": true
         },
         "phone_number": {
            "type": "String"
         },
         "phone_number_verified": {
            "type": "Boolean"
         }
      }
   }
}
```

For an example policy that will validate against this schema, see Reflects OIDC ID token attributes.

Mapping access tokens

Verified Permissions processes access-token claims other than the groups claim as attributes of the action, or *context attributes*. Along with group membership, the access tokens from your IdP might contain information about API access. Access tokens are useful in authorization models that use role-based access control (RBAC). Authorization models that rely on access-token claims other than group membership require additional effort in schema configuration.

Most access tokens from external OIDC providers align closely with Amazon Cognito access tokens. An OIDC access token is mapped to a context object when passed to Verified Permissions. Attributes of the access token can be referenced using `context.token.`*`attribute_name`*. The following example OIDC access token includes example base claims.

```
{
    "sub": "91eb4550-9091-708c-a7a6-9758ef8b6b1e",
    "groups": [
        "Store-Owner-Role",
        "Customer"
    ],
    "iss": "https://auth.example.com",
    "client_id": "1example23456789",
    "aud": "https://myapplication.example.com"
    "scope": "MyAPI-Read",
    "exp": 1688096566,
    "iat": 1688092966,
    "jti": "a1b2c3d4-e5f6-a1b2-c3d4-TOKEN2222222",
    "username": "alice"
```

```
}
```

The following example shows how to reflect the attributes from the example access token in your Verified Permissions schema. For more information about editing your schema, see Editing policy store schemas.

```
{
   "MyApplication": {
      "actions": {
         "Read": {
            "appliesTo": {
               "context": {
                  "type": "ReusedContext"
               },
               "resourceTypes": [
                  "Application"
               ],
               "principalTypes": [
                  "User"
               ]
            }
         }
      },
      ...
      ...
      "commonTypes": {
         "ReusedContext": {
            "attributes": {
               "token": {
                  "type": "Record",
                  "attributes": {
                     "scope": {
                        "type": "Set",
                        "element": {
                           "type": "String"
                        }
                     },
                     "client_id": {
                        "type": "String"
                     }
                  }
               },
```

```
            "type": "Record"
         }
      }
   }
}
```

For an example policy that will validate against this schema, see Reflects OIDC access token attributes.

Things to know about schema mapping

Attribute mapping differs between token types

In access token authorization, Verified Permissions maps claims to context. In ID token authorization, Verified Permissions maps claims to principal attributes. For policy stores that you create in the Verified Permissions console, only **empty** and **sample** policy stores leave you with no identity source and require you to populate your schema with user pool attributes for ID token authorization. Access token authorization is based on role-based access control (RBAC) with group-membership claims and doesn't automatically map other claims to the policy store schema.

Identity source attributes aren't required

When you create an identity source in the Verified Permissions console, no attributes are marked as required. This prevents missing claims from causing validation errors in authorization requests. You can set attributes to required as needed, but they must be present in all authorization requests.

RBAC doesn't require attributes in schema

Schemas for identity sources depend on the entity associations that you make when you add your identity source. An identity source maps one claim to a user entity type, and one claim to a group entity type. These entity mappings are the core of an identity-source configuration. With this minimum information, you can write policies that perform authorization actions for specific users and specific groups that users might be members of, in a role-based access control (RBAC) model. The addition of token claims to the schema extends the authorization scope of your policy store. User attributes from ID tokens have information about users that can contribute to attribute-based access control (ABAC) authorization. Context attributes from access tokens have information like OAuth 2.0 scopes that can contribute additional access-control information from your provider, but require additional schema modifications.

The **Set up with API Gateway and an identity provider** and **Guided setup** options in the Verified Permissions console assign ID token claims to the schema. This isn't the case for access token

claims. To add non-group access-token claims to your schema, you must edit your schema in JSON mode and add commonTypes attributes. For more information, see Mapping access tokens.

OIDC groups claim supports multiple formats

When you add an OIDC provider, you can choose the name of the groups claim in ID or access tokens that you want to map to a user's group membership in your policy store. Verified permissions recognizes groups claims in the following formats:

1. String without spaces: `"groups": "MyGroup"`
2. Space-delimited list: `"groups": "MyGroup1 MyGroup2 MyGroup3"`. Each string is a group.
3. JSON (comma-delimited) list: `"groups": ["MyGroup1", "MyGroup2", "MyGroup3"]`

> **Note**
>
> Verified Permissions interprets each string in a space-separated groups claim as a separate group. To interpret a group name with a space character as a single group, replace or remove the space in the claim. For example, format a group named `My Group` as `MyGroup`.

Choose a token type

The way that your policy store works with your identity source depends on a key decision in identity-source configuration: whether you will process ID or access tokens. With an OIDC provider, you must choose a token type when you add the identity source. You can choose ID or access token, and your choice excludes the unchosen token type from being processed in your policy store. Especially if you wish to benefit from the automatic mapping of ID token claims to attributes in the Verified Permissions console, decide early about the token type that you want to process before you create your identity source. Changing the token type requires significant effort to refactor your policies and schema. The following topics describe the use of ID and access tokens with policy stores.

Cedar parser requires brackets for some characters

Policies typically reference schema attributes in a format like `principal.username`. In the case of most non-alphanumeric characters like `:`, `.`, or `/` that might appear in token claim names, Verified Permissions can't parse a condition value like `principal.cognito:username` or `context.ip-address`. You must instead format these conditions with bracket notation in

the format `principal["cognito:username"]` or `context["ip-address"]`, respectively. The underscore character _ is a valid character in claim names, and the only non-alphanumeric exception to this requirement.

A partial example schema for a principal attribute of this type looks like the following:

```
"User": {
   "shape": {
      "type": "Record",
      "attributes": {
         "cognito:username": {
            "type": "String",
            "required": true
         },
         "custom:employmentStoreCode": {
            "type": "String",
            "required": true,
         },
         "email": {
            "type": "String",
            "required": false
         }
      }
   }
}
```

A partial example schema for a context attribute of this type looks like the following:

```
"GetOrder": {
   "memberOf": [],
   "appliesTo": {
      "resourceTypes": [
         "Order"
      ],
      "context": {
         "type": "Record",
         "attributes": {
            "ip-address": {
               "required": false,
               "type": "String"
            }
   }
   },
```

```
      "principalTypes": [
         "User"
      ]
   }
}
```

For an example policy that will validate against this schema, see Uses bracket notation to reference token attributes.

Client and audience validation for OIDC providers

When you add an identity source to a policy store, Verified Permissions has configuration options that verify that ID and access tokens are being used as intended. This validation happens in the processing of `IsAuthorizedWithToken` and `BatchIsAuthorizedWithToken` API requests. The behavior differs between ID and access tokens, and between Amazon Cognito and OIDC identity sources. With Amazon Cognito user pools providers, Verified Permissions can validate the client ID in both ID and access tokens. With OIDC providers, Verified Permissions can validate the client ID in ID tokens, and the audience in access tokens.

A *client ID* is an identifier associated with the identity provider instance that your application uses, for example `1example23456789`. An *audience* is a URL path associated with the intended *relying party*, or destination, of the access token, for example `https://mytoken.example.com`. When using access tokens, the `aud` claim is always associated with the audience.

OIDC ID tokens have an `aud` claim that contains client IDs, such as `1example23456789`.

OIDC Access tokens have an `aud` claim that contains the audience URL for the token, such as `https://myapplication.example.com`, and a `client_id` claim that contains client IDs, such as `1example23456789`.

When setting up your policy store, enter one or more values for **Audience validation** that your policy store with use to validate the audience of a token.

- **ID tokens** – Verified Permissions validates the client ID by checking that at least one member of the client IDs in the `aud` claim matches an audience validation value.
- **Access tokens** – Verified Permissions validates the audience by checking that the URL in the `aud` claim matches an audience validation value. If no `aud` claim exists, the audience can be validated using the `cid` or `client_id` claims. Check with your identity provider for the correct audience claim and format.

Client-side authorization for JWTs

You might want to process JSON web tokens in your application and pass their claims to Verified Permissions without using a policy store identity source. You can extract your entity attributes from a JSON Web Token (JWT) and parse it into Verified Permissions.

This example shows how you might call Verified Permissions from an application using a JWT.[1]

```
async function authorizeUsingJwtToken(jwtToken) {

    const payload = await verifier.verify(jwtToken);

    let principalEntity = {
        entityType: "PhotoFlash::User", // the application needs to fill in the
 relevant user type
        entityId: payload["sub"], // the application need to use the claim that
 represents the user-id
    };
    let resourceEntity = {
        entityType: "PhotoFlash::Photo", //the application needs to fill in the
 relevant resource type
        entityId: "jane_photo_123.jpg", // the application needs to fill in the
 relevant resource id
    };
    let action = {
        actionType: "PhotoFlash::Action", //the application needs to fill in the
 relevant action id
        actionId: "GetPhoto", //the application needs to fill in the relevant action
 type
    };
    let entities = {
        entityList: [],
    };
    entities.entityList.push(...getUserEntitiesFromToken(payload));
    let policyStoreId = "PSEXAMPLEabcdefg111111"; // set your own policy store id

    const authResult = await client
        .isAuthorized({
        policyStoreId: policyStoreId,
        principal: principalEntity,
        resource: resourceEntity,
        action: action,
        entities,
```

```
        })
        .promise();

    return authResult;

}

function getUserEntitiesFromToken(payload) {
  let attributes = {};
  let claimsNotPassedInEntities = ['aud', 'sub', 'exp', 'jti', 'iss'];
  Object.entries(payload).forEach(([key, value]) => {
    if (claimsNotPassedInEntities.includes(key)) {
        return;
    }
    if (Array.isArray(value)) {
      var attibuteItem = [];
      value.forEach((item) => {
        attibuteItem.push({
          string: item,
        });
      });
      attributes[key] = {
        set: attibuteItem,
      };
    } else if (typeof value === 'string') {
      attributes[key] = {
        string: value,
      }
    } else if (typeof value === 'bigint' || typeof value ==='number') {
        attributes[key] = {
            long: value,
          }
    } else if (typeof value === 'boolean') {
        attributes[key] = {
            boolean: value,
       }
    }

  });

  let entityItem = {
    attributes: attributes,
    identifier: {
      entityType: "PhotoFlash::User",
```

```
      entityId: payload["sub"], // the application needs to use the claim that
 represents the user-id
    }
  };
  return [entityItem];
}
```

[1] This code example uses the aws-jwt-verify library for verifying JWTs signed by OIDC-compatible IdPs.

Integrations for Amazon Verified Permissions

Amazon Verified Permissions integrations help you implement fine-grained authorization in your applications while minimizing code and following framework-specific best practices. These integrations provide middleware components and utilities that seamlessly connect your application with Verified Permissions.

With integrations, you can:

- Implement authorization in minutes
- Follow framework-specific patterns and conventions
- Reduce maintenance overhead
- Minimize potential security implementation errors
- Focus on business logic rather than authorization code

When added to your application, integrations do the following:

1. Intercept incoming requests through framework-specific middleware
2. Extract relevant authorization context from requests
3. Determine authorization decisions using Verified Permissions
4. Enforce access control based on authorization results

Verified Permissions currently supports the following frameworks:

- Express.js for Node.js applications

Integrating Express with Amazon Verified Permissions

The Verified Permissions Express integration provides a middleware-based approach to implementing authorization in your Express.js applications. With this integration, you can protect your API endpoints using fine-grained authorization policies without modifying your existing route handlers. The integration handles authorization checks automatically by intercepting requests, evaluating them against your defined policies, and ensuring that only authorized users can access protected resources.

This topic walks you through setting up the Express integration, from creating a policy store to implementing and testing the authorization middleware. By following these steps, you can add robust authorization controls to your Express application with minimal code changes.

The following GitHub repos are referenced throughout this topic:

- cedar-policy/authorization-for-expressjs - The Cedar authorization middleware for Express.js
- verifiedpermissions/authorization-clients-js - The Verified Permissions authorization clients for JavaScript
- verifiedpermissions/examples/express-petstore - Example implementation using the Express.js middleware

Prerequisites

Before you implement the Express integration, ensure you have:

- An AWS account with access to Verified Permissions
- Node.js and npm installed
- An Express.js application
- An OpenID Connect (OIDC) identity provider (such as Amazon Cognito)
- AWS CLI configured with appropriate permissions

Setting up the integration

Step 1: Create a policy store

Create a policy store using the AWS CLI:

```
aws verifiedpermissions create-policy-store --validation-settings "mode=STRICT"
```

> **Note**
>
> Save the policy store ID returned in the response for use in subsequent steps.

Step 2: Install dependencies

Install the required packages in your Express application:

```
npm i --save @verifiedpermissions/authorization-clients-js
npm i --save @cedar-policy/authorization-for-expressjs
```

Configuring authorization

Step 1: Generate and upload Cedar schema

A schema defines the authorization model for an application, including the entities types in the application and the actions users are allowed to take. We recommend defining a namespace for your schema. In this example, we use `YourNamespace`. You attach your schema to your Verified Permissions policy stores, and when policies are added or modified, the service automatically validates the policies against the schema.

The `@cedar-policy/authorization-for-expressjs` package can analyze the OpenAPI specifications of your application and generate a Cedar schema. Specifically, the paths object is required in your specification.

If you don't have an OpenAPI specification, you can follow the quick instructions of the express-openapi-generator package to generate an OpenAPI specification.

Generate a schema from your OpenAPI specification:

```
npx @cedar-policy/authorization-for-expressjs generate-schema --api-spec schemas/
openapi.json --namespace YourNamespace --mapping-type SimpleRest
```

Next, format the Cedar schema for use with the AWS CLI. For more information about the specific format required, see *Policy store schema*. If you need help formatting the schema, there's a script called `prepare-cedar-schema.sh` in the verifiedpermissions/examples GitHub repo. The following is an example call to that script that outputs the Verified Permissions formatted schema in the `v2.cedarschema.forAVP.json` file.

```
./scripts/prepare-cedar-schema.sh v2.cedarschema.json v2.cedarschema.forAVP.json
```

Upload the formatted schema to your policy store, replacing `policy-store-id` with your policy store ID:

```
aws verifiedpermissions put-schema \
  --definition file://v2.cedarschema.forAVP.json \
  --policy-store-id policy-store-id
```

Step 2: Create authorization policies

If no policies are configured, Cedar denies all authorization requests. The Express framework integration helps bootstrap this process by generating example policies based on the previously generated schema.

When using this integration in your production applications, we recommend creating new policies using infrastructure as a code (IaaC) tools. For more information, see *Working with AWS CloudFormation*.

Generate sample Cedar policies:

```
npx @cedar-policy/authorization-for-expressjs generate-policies --schema
 v2.cedarschema.json
```

This will generate sample policies in the `/policies` directory. You can then customize these policies based on your use cases. For example:

```
// Defines permitted administrator user group actions
permit (
    principal in YourNamespace::UserGroup::"<userPoolId>|administrator",
    action,
    resource
);

// Defines permitted employee user group actions
permit (
    principal in YourNamespace::UserGroup::"<userPoolId>|employee",
    action in
        [YourNamespace::Action::"GET /resources",
         YourNamespace::Action::"POST /resources",
         YourNamespace::Action::"GET /resources/{resourceId}",
         YourNamespace::Action::"PUT /resources/{resourceId}"],
    resource
);
```

Format the policies for use with the AWS CLI. Fore more information about the required format, see create-policy in the *AWS CLI reference*. If you need help formatting the policies, there's a script called `convert_cedar_policies.sh` in the verifiedpermissions/examples GitHub repo. The following is a call to that script:

```
./scripts/convert_cedar_policies.sh
```

Upload the formatted policies to Verified Permissions, replacing `policy_1.json` with the path and name of your policy file and `policy-store-id` with your policy store ID:

```
aws verifiedpermissions create-policy \
  --definition file://policies/json/policy_1.json \
  --policy-store-id policy-store-id
```

Step 3: Connect an identity provider

By default, the Verified Permissions authorizer middleware reads a JSON Web Token (JWT) provided within the authorization header of the API request to get user information. Verified Permissions can validate the token in addition to performing authorization policy evaluation.

Create an identity source configuration file named `identity-source-configuration.txt` that looks like the following with your `userPoolArn` and `clientId`:

```
{
    "cognitoUserPoolConfiguration": {
        "userPoolArn": "arn:aws:cognito-idp:region:account:userpool/pool-id",
        "clientIds": ["client-id"],
        "groupConfiguration": {
            "groupEntityType": "YourNamespace::UserGroup"
        }
    }
}
```

Create the identity source by running the following AWS CLI command, replacing `policy-store-id` with your policy store ID:

```
aws verifiedpermissions create-identity-source \
  --configuration file://identity-source-configuration.txt \
  --policy-store-id policy-store-id \
  --principal-entity-type YourNamespace::User
```

Implementing the authorization middleware

Update your Express application to include the authorization middleware. In this example we're using identity tokens, but you can also use access tokens. For more information, see authorization-for-expressjs on GitHub.

```
const { ExpressAuthorizationMiddleware } = require('@cedar-policy/authorization-for-
expressjs');

const { AVPAuthorizationEngine } = require('@verifiedpermissions/authorization-
clients');

const avpAuthorizationEngine = new AVPAuthorizationEngine({
    policyStoreId: 'policy-store-id',
    callType: 'identityToken'
});

const expressAuthorization = new ExpressAuthorizationMiddleware({
    schema: {
        type: 'jsonString',
        schema: fs.readFileSync(path.join(__dirname, '../v4.cedarschema.json'),
 'utf8'),
    },
    authorizationEngine: avpAuthorizationEngine,
    principalConfiguration: { type: 'identityToken' },
    skippedEndpoints: [],
    logger: {
        debug: (s) => console.log(s),
        log: (s) => console.log(s),
    }
});

// Add the middleware to your Express application
app.use(expressAuthorization.middleware);
```

Testing the integration

You can test your authorization implementation by making requests to your API endpoints with different user tokens. The authorization middleware will automatically evaluate each request against your defined policies.

For example, if you've set up different user groups with different permissions:

- Administrators: Full access to all resources and management functions
- Employees: Can view, create, and update resources
- Customers: Can only view resources

You can validate that the permissions policies are working as expected by signing in with different users and attempting various operations. In the terminal for the Express application, you can see log output that provides additional details about the authorization decisions.

Troubleshooting

If you have authorization failures, try the following:

- Verify your policy store ID is correct
- Ensure your identity source is properly configured
- Check that your policies are correctly formatted
- Validate that your JWT tokens are valid

Next steps

After implementing the basic integration, consider:

- Implementing custom mappers for specific authorization scenarios
- Setting up monitoring and logging for authorization decisions
- Creating additional policies for different user roles

Implementing authorization in Amazon Verified Permissions

After you build your policy store, policies, templates, schema, and authorization model, you're ready to start authorizing requests using Amazon Verified Permissions. To implement Verified Permissions authorization, you must combine configuration of authorization policies in AWS with integration in an application. To integrate Verified Permissions with your application, add an AWS SDK and implement the methods that invoke the Verified Permissions API and generate authorization decisions against your policy store.

Authorization with Verified Permissions is useful for *UX permissions* and *API permissions* in your applications.

UX permissions

Control user access to your application UX. You can permit a user to view only the exact forms, buttons, graphics and other resources that they need to access. For example, when a user signs in, you might want to determine whether a "Transfer funds" button is visible in their account. You can also control actions that a user can take. For example, in same banking app you might want to determine whether your user is permitted to change the category of a transaction.

API permissions

Control user access to data. Applications are often part of a distributed system and bring in information from external APIs. In the example of the banking app where Verified Permissions has permitted the display of a "Transfer funds" button, a more complex authorization decision must be made when your user initiates a transfer. Verified Permissions can authorize the API request that lists the destination accounts that are eligible transfer targets, and then the request to push the transfer to the other account.

The examples that illustrate this content come from a sample policy store. To follow along, create the **DigitalPetStore** sample policy store in your testing environment.

For an end to end sample application that implements UX permissions using batch authorization, see Use Amazon Verified Permissions for fine-grained authorization at scale on the *AWS Security Blog*.

Topics

- Available API operations for authorization
- Testing your authorization model
- Integrating your authorization models with applications

Available API operations for authorization

The Verified Permissions API has the following authorization operations.

IsAuthorized

The `IsAuthorized` API operation is the entry point to authorization requests with Verified Permissions. You must submit principal, action, resource, context, and entities elements. Verified Permissions evaluates your request against all policies in the requested policy store that apply to the entities in the request.

IsAuthorizedWithToken

The `IsAuthorizedWithToken` operation generates an authorization request from user data in JSON web tokens (JWTs). Verified Permissions works directly with OIDC providers like Amazon Cognito as an identity source in your policy store. Verified Permissions populates all attributes to the principal in your request from the claims in users' ID or access tokens. You can authorize actions and resources from user attributes or group membership in an identity source.

You can't include information about group or user principal types in an `IsAuthorizedWithToken` request. You must populate all principal data to the JWT that you provide.

BatchIsAuthorized

The `BatchIsAuthorized` operation processes multiple authorization decisions for a single principal or resource in a single API request. This operation groups requests into a single batch operation that minimizes quota usage and returns authorization decisions for each of up to 30 complex nested actions. With batch authorization for a single resource, you can filter the actions that a user can take on a resource. With batch authorization for a single principal, you can filter for the resources that a user can take action on.

BatchIsAuthorizedWithToken

The `BatchIsAuthorizedWithToken` operation processes multiple authorization decisions for a single principal in one API request. The principal is provided by your policy store identity source in an ID or access token. This operation groups requests into a single batch operation

that minimizes quota usage and returns authorization decisions for each of up to 30 requests for actions and resources. In your policies, you can authorize their access from their attributes or their group membership in a user directory.

Like with `IsAuthorizedWithToken`, you can't include information about group or user principal types in a `BatchIsAuthorizedWithToken` request. You must populate all principal data to the JWT that you provide.

Testing your authorization model

To understand the effect of Amazon Verified Permissions authorization decision when you deploy your application, you can evaluate your policies as you develop them with the Using the Amazon Verified Permissions test bench and with HTTPS REST API requests to Verified Permissions. The test bench is a tool in the AWS Management Console to evaluate authorization requests and responses in your policy store.

The Verified Permissions REST API is the next step in your development as you move from a conceptual understanding to application design. The Verified Permissions API accepts authorization requests with IsAuthorized, IsAuthorizedWithToken, and BatchIsAuthorized as signed AWS API requests to Regional service endpoints. To test your authorization model, you can generate requests with any API client and verify that your policies are returning authorization decisions as expected.

For example, you can test `IsAuthorized` in a sample policy store with the following procedure.

Test bench

1. Open the Verified Permissions console at Verified Permissions console. Create a policy store from the **Sample policy store** with the name **DigitalPetStore**.
2. Select **Test bench** in your new policy store.
3. Populate your test bench request from IsAuthorized in the Verified Permissions API reference. The following details replicate the conditions in **Example 4** that references the **DigitalPetStore** sample.

    a. Set Alice as the principal. For **Principal taking action**, choose `DigitalPetStore::User` and enter `Alice`.

    b. Set Alice's role as customer. Choose **Add a parent**, choose `DigitalPetStore::Role`, and enter Customer.

c. Set the resource as order "1234." For **Resource that the principal is acting on**, choose `DigitalPetStore::Order` and enter `1234`.

d. The `DigitalPetStore::Order` resource requires an `owner` attribute. Set Alice as the owner of the order. Choose `DigitalPetStore::User` and enter `Alice`

e. Alice requested to view the order. For **Action that principal is taking**, choose `DigitalPetStore::Action::"GetOrder"`.

4. Choose **Run authorization request**. In an unmodified policy store, this request results in an `ALLOW` decision. Note the **Satisfied policy** that returned the decision.
5. Choose **Policies** from the left navigation bar. Review the static policy with the description **Customer Role - Get Order**.
6. Observe that Verified Permissions allowed the request because the principal was in a customer role and was the owner of the resource.

REST API

1. Open the Verified Permissions console at Verified Permissions console. Create a policy store from the **Sample policy store** with the name **DigitalPetStore**.
2. Note the **Policy store ID** of your new policy store.
3. From IsAuthorized in the Verified Permissions API reference, copy the request body of **Example 4** that references the **DigitalPetStore** sample.
4. Open your API client and create a request to the Regional service endpoint for your policy store. Populate the headers as shown in the example.
5. Paste in the sample request body and change the value of `policyStoreId` to the policy store ID you noted earlier.
6. Submit the request and review the results. In a default **DigitalPetStore** policy store, this request returns an `ALLOW` decision.

You can make changes to policies, schema, and requests in your test environment to change the outcomes and produce more complex decisions.

1. Change the request in a way that changes the decision from Verified Permissions. For example, change Alice's role to `Employee` or change the `owner` attribute of order 1234 to Bob.

2. Change policies in ways that affect authorization decisions. For example, modify the policy with the description **Customer Role - Get Order** to remove the condition that the `User` must be the owner of the `Resource` and modify the request so that Bob wants to view the order.
3. Change the schema to allow policies to make a more complex decision. Update the request entities so that Alice can satisfy the new requirements. For example, edit the schema to allow `User` to be a member of `ActiveUsers` or `InactiveUsers`. Update the policy so that only active users can view their own orders. Update the request entities so that Alice is an active or inactive user.

Integrating your authorization models with applications

To implement Amazon Verified Permissions in your application, you must define the policies and schema that you want your app to enforce. With your authorization model in place and tested, your next step is to start generating API requests from the point of enforcement. To do this, you must set up application logic to collect user data and populate it to authorization requests.

How an app authorizes requests with Verified Permissions

1. Gather information about the current user. Typically, a user's details are provided in the details of an authenticated session, like a JWT or web session cookie. This user data might originate from an Amazon Cognito identity source linked to your policy store or from another OpenID Connect (OIDC) provider.
2. Gather information about the resource that a user wants to access. Typically, your application will receive information about the resource when a user makes a selection that requires your app to load a new asset.
3. Determine the action that your user wants to take.
4. Generate an authorization request to Verified Permissions with the principal, action, resource, and entities for your user's attempted operation.Verified Permissions evaluates the request against the policies in your policy store and returns an authorization decision.
5. Your application reads the allow or deny response from Verified Permissions and enforces the decision on the user's request.

Verified Permissions API operations are built into AWS SDKs. To include Verified Permissions in an app, integrate the AWS SDK for your chosen language into the app package.

To learn more and download AWS SDKs, see Tools for Amazon Web Services.

The following are links to documentation for Verified Permissions resources in various AWS SDKs.

- AWS SDK for .NET
- AWS SDK for C++
- AWS SDK for Go
- AWS SDK for Java
- AWS SDK for JavaScript
- AWS SDK for PHP
- AWS SDK for Python (Boto)
- AWS SDK for Ruby
- AWS SDK for Rust

The following AWS SDK for JavaScript example for `IsAuthorized` originates from Simplify fine-grained authorization with Amazon Verified Permissions and Amazon Cognito.

```
const authResult = await avp.isAuthorized({
    principal: 'User::"alice"',
    action: 'Action::"view"',
    resource: 'Photo::"VacationPhoto94.jpg"',
    // whenever our policy references attributes of the entity,
    // isAuthorized needs an entity argument that provides
    // those attributes
    entities: {
       entityList: [
         {
            "identifier": {
                "entityType": "User",
                "entityId": "alice"
            },
            "attributes": {
                "location": {
                    "String": "USA"
                }
            }
         }
       ]
    }
});
```

More developer resources

- Amazon Verified Permissions workshop
- Amazon Verified Permissions - Resources
- Implement custom authorization policy provider for ASP.NET Core apps using Amazon Verified Permissions
- Build an entitlement service for business applications using Amazon Verified Permissions
- Simplify fine-grained authorization with Amazon Verified Permissions and Amazon Cognito

Making API requests

Query requests for the Amazon Verified Permissions are HTTP or HTTPS requests that use the POST method.

Verified Permissions endpoints

An *endpoint* is a URL that serves as an entry point for a web service. You can select an appropriate AWS Region endpoint when you make your requests to reduce latency. For information about the endpoints used by Verified Permissions, see Amazon Verified Permissions in the *Amazon Web Services General Reference*.

Request parameters

Amazon Verified Permissions uses a JSON-based protocol. You pass request parameters as JSON in the body of your HTTP request. Lists are represented as JSON arrays. For more information, see Common Parameters in the *Amazon Verified Permissions API Reference*.

Request identifiers

In every response, AWS returns a request ID via the `x-amzn-RequestId` HTTP response header. This string is a unique identifier that AWS assigns to provide tracking information. Although the request ID is included in every response, it isn't listed on the individual API documentation pages to improve readability and reduce redundancy.

Query API authentication

You send query requests over HTTPS. You must include a signature in every query request. For more information about creating and including a signature, see Signing AWS API Requests in the *Amazon Web Services General Reference*.

Available libraries

AWS provides libraries, sample code, tutorials, and other resources for software developers who prefer to build applications using language-specific APIs instead of the command-line tools and

Query API. These libraries provide basic functions (not included in the APIs), such as request authentication, request retries, and error handling so that it's easier to get started. Verified Permissions libraries and resources are available for the following languages and platforms:

- AWS SDK for Go
- AWS SDK for Java 2.x
- AWS SDK for Java 1.x
- AWS SDK for JavaScript
- AWS SDK for .NET
- AWS SDK for PHP
- AWS SDK for Python (Boto)
- AWS SDK for Ruby

For more information about libraries and sample code in all languages, see Sample Code & Libraries.

Making API requests using the POST method

If you don't use one of the AWS SDKs, you can make Verified Permissions requests over HTTPS using the POST request method. The POST method requires that you specify the operation in the header of the request and provide the data for the operation in JSON format in the body of the request.

Header name	Header value
`Host`	The Amazon Verified Permissions endpoint. For example: `verifiedpermissions.us-east-1.amazonaws.com`
`X-Amz-Date`	You must provide the timestamp in either the HTTP Date header or the AWS *x-amz-date* header. Some HTTP client libraries don't let you set the Date header. When an *x-amz-date* header is present, the system ignores any Date header during the request authentication. The *x-amz-date* header must be specified in ISO 8601 basic format. For example: `20130315T092054Z`

Header name	Header value
`Authorization`	The set of authorization parameters that AWS uses to ensure the validity and authenticity of the request. For more information about constructing this header, see Signature Version 4 Signing Process in the *Amazon Web Services General Reference*.
`X-Amz-Target`	Specifies the Verified Permissions operation that you want to perform. `VerifiedPermissions.` *`API_Name`* For example, to call the `CreatePolicy` operation, use the following target value. `VerifiedPermissions.CreatePolicy`
`Content-Type`	Specifies the input format. Use the following value. `application/x-amz-json-1.0`
`Accept`	Specifies the response format. Use the following value. `application/x-amz-json-1.0`
`Content-Length`	Size of the payload in bytes.
`Content-Encoding`	Specifies the encoding format of the input and output. Use the following value. `amz-1.0`

The following is an example header for an HTTP request to return a list of all policies in the specified policy store in the AWS account where the `Principal` references a `User` named `alice`. In this example, the `Authorization` line is word-wrapped here for easier reading. Don't word wrap it in your actual request.

```
POST / HTTP/1.1
Host: verifiedpermissions.us-east-1.amazonaws.com
X-Amz-Date: 20230101T200059Z
Accept-Encoding: identity
```

```
Content-Type: application/x-amz-json-1.0
X-Amz-Target: VerifiedPermissions.ListPolicies
User-Agent: <UserAgentString>
Authorization: AWS4-HMAC-SHA256 Credential=<Credential>, SignedHeaders=<Headers>,
 Signature=<Signature>
Content-Length: <PayloadSizeBytes>

{
    "PolicyStoreId": "PSExAmPLE222222222",
    "Filter": {
        "Principal": {
            "Id": {
                "EntityType": "User",
                "EntityId": "alice"
            }
        }
    }
}
```

Security in Amazon Verified Permissions

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from data centers and network architectures that are 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 Amazon Verified Permissions, 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 Verified Permissions. The following topics show you how to configure Verified Permissions to meet your security and compliance objectives. You also learn how to use other AWS services that help you to monitor and secure your Verified Permissions resources.

Topics

Data protection in Amazon Verified Permissions

The AWS shared responsibility model applies to data protection in Amazon Verified Permissions. As described in this model, AWS is responsible for protecting the global infrastructure that runs all of the AWS Cloud. You are responsible for maintaining control over your content that is hosted on this infrastructure. This content includes the security configuration and management tasks for the

AWS services that you use. For more information about data privacy, see the Data Privacy FAQ. For information about data protection in Europe, see the AWS Shared Responsibility Model and GDPR blog post on the *AWS Security Blog*.

- For data protection purposes, we recommend that you protect AWS account credentials and set up individual users with AWS IAM Identity Center or AWS Identity and Access Management (IAM). That way, each user is given only the permissions necessary to fulfill their job duties.
- We recommend that you secure your data in the following ways:
  - Use multi-factor authentication (MFA) with each account.
  - Use SSL/TLS to communicate with AWS resources. We require TLS 1.2.
  - Set up API and user activity logging with AWS CloudTrail.
  - Use AWS encryption solutions, along with all default security controls within AWS services.
  - Use advanced managed security services such as Amazon Macie, which assists in discovering and securing sensitive data that is stored in Amazon S3.
  - If you require FIPS 140-2 validated cryptographic modules when accessing AWS through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see Federal Information Processing Standard (FIPS) 140-2.
- We strongly recommend that you never put confidential or sensitive information, such as your customers' email addresses, into tags or free-form text fields such as a **Name** field. This includes when you work with Verified Permissions or other AWS services using the console, API, AWS CLI, or AWS SDKs. Any data that you enter into tags or free-form text fields used for names may be used for billing or diagnostic logs. If you provide a URL to an external server, we strongly recommend that you do not include credentials information in the URL to validate your request to that server.
- Your action names should not include any sensitive information.
- We also strongly recommend that you always use unique, non-mutable, and non-reusable identifiers for your entities (resources and principals). In a test environment, you might choose to use simple entity identifiers, such as `jane` or bob for the name of an entity of type `User`. However, in a production system, it's critical for security reasons that you use unique values that can't be reused. We recommend that you use values like universally unique identifiers (UUIDs). For example, consider the user `jane` who leaves the company. Later, you let someone else use the name `jane`. That new user gets access automatically to everything granted by policies that still reference `User::"jane"`. Verified Permissions and Cedar can't distinguish between the new user and the previous user.

This guidance applies to both principal and resource identifiers. Always use identifiers that are guaranteed unique and never reused to ensure that you don't grant access unintentionally because of the presence of an old identifier in a policy.

- Ensure that the strings that you provide to define `Long` and `Decimal` values are within the valid range of each type. Also, ensure that your use of any arithmetic operators don't result in a value outside of the valid range. If the range is exceeded, the operation results in an overflow exception. A policy that results in an error is ignored, meaning that a Permit policy might unexpectedly fail to allow access, or a Forbid policy might unexpectedly fail to block access.

Data encryption

Amazon Verified Permissions automatically encrypts all customer data such as policies with an AWS managed key. Amazon Verified Permissions also allows for customers to utilize a customer managed key to encrypt their data.

For detailed information about using customer managed keys for encryption, see the section called "Customer managed keys".

Encrypting Resources in Amazon Verified Permissions

Amazon Verified Permissions provides encryption by default to protect sensitive customer data at rest using AWS owned encryption keys. As an extra layer of protection, Amazon Verified Permissions allows you to encrypt your policy stores using AWS Key Management Service (AWS KMS) customer managed keys (CMK). This functionality ensures protection of sensitive data via encryption at rest, which helps you:

- Reduce the operational burden on your application's end to protect sensitive data
- Maintain control over who can see details of your authorization policies via your own AWS KMS customer managed keys
- Build security-sensitive applications that meet strict encryption compliance and regulatory requirements

The following sections explain how to configure encryption for new policy stores and managing your encryption keys.

AWS KMS Key Types for Amazon Verified Permissions

Amazon Verified Permissions integrates with AWS KMS to manage encryption keys used for encrypting/decrypting customer data. To learn more about key types and states, see AWS Key Management Service concepts in the *AWS KMS Developer Guide*. When you create a new policy store, you can choose from the following AWS KMS key types to encrypt your data:

AWS Owned Key

The default encryption type. Amazon Verified Permissions owns the key at no additional charge to you and encrypts resource data at rest upon creation. No additional configuration is required in your code or applications to encrypt/decrypt your data using the key owned by Verified Permissions.

Customer Managed Key

You create, own, and manage the key in your AWS account. You have full control over the AWS KMS key. AWS KMS charges apply for customer managed keys. For more information, see the AWS KMS Pricing page. For more information about key types, see Customer managed keys in the *AWS KMS Developer Guide*.

When you specify a customer managed key for encryption for top-level resources (i.e. policy store), Verified Permissions encrypts the resource, as well as its child resources, with that key. To encrypt a policy store using a customer managed key, you need to grant access to Verified Permissions in your key policy. A key policy is a resource-based policy that you attach to your customer managed key to control access to it. See the section called "Authorizing use of your AWS KMS key for Amazon Verified Permissions" for more details.

In addition, to create an encrypted policy store with a customer managed key, or to make API calls to a policy store encrypted by a customer managed key, the IAM user or role which makes the call must also have access to the key. If Verified Permissions is unable to access the key, any authorization decisions that involve resources encrypted by that key may be stale or inaccurate. When you do not have access to the key, you will not be able to read/update/delete resources encrypted by that key, and any create calls to utilize the key for encryption will fail.

> **Note**
>
> Verified Permissions encryption at rest is available in all AWS Regions where Verified Permissions is available.

> **Important**
>
> Once a customer managed key has been used to encrypt a policy store, you CANNOT update the resource to use a different key for encryption or remove the key from that policy store.

Using AWS KMS and data keys with Amazon Verified Permissions

The Amazon Verified Permissions encryption at rest feature uses an AWS KMS key and a hierarchy of data keys to protect your resource data.

> **Note**
>
> Amazon Verified Permissions supports only symmetric AWS KMS keys. You can't use an asymmetric AWS KMS key to encrypt your Amazon Verified Permissions resources.

Using AWS Owned Keys

Amazon Verified Permissions encrypts all resources by default with AWS owned keys. These keys are free to use and rotate annually to protect your account resources. You don't need to view, manage, use, or audit these keys, so there's no action required for data protection. For more information about AWS owned keys, see AWS owned keys in the *AWS KMS Developer Guide*.

Using Customer Managed Keys

Selecting a customer managed key for encryption provides the following benefits:

- You create and manage the AWS KMS key, including setting the key policies and IAM policies to control access to the AWS KMS key. You can enable and disable the AWS KMS key, enable and disable automatic key rotation, and delete the AWS KMS key when it is no longer in use.
- You can use a customer managed key with imported key material or a customer managed key in a custom key store that you own and manage.
- You can audit the encryption and decryption of your Verified Permissions resources by examining the Amazon Verified Permissions API calls to AWS KMS in AWS CloudTrail logs.

For Amazon Verified Permissions to use your customer managed keys for encryption/decryption, you will need to add specific key policies to allow Amazon Verified Permissions to encrypt/decrypt resources on your behalf.

Authorizing use of your AWS KMS key for Amazon Verified Permissions

At a minimum, Amazon Verified Permissions requires the following permissions on a customer managed key:

- `kms:Encrypt`
- `kms:GenerateDataKeyWithoutPlaintext`
- `kms:DescribeKey`
- `kms:ReEncryptTo`
- `kms:ReEncryptFrom`
- `kms:Decrypt`

An example key policy can be seen below:

```
{
    "Sid": "Enable AVP to use the KMS key for encrypting project J.A.K. policy
 resources",
    "Effect": "Allow",
    "Principal": {
        "Service": "verifiedpermissions.amazonaws.com"
    },
    "Action": [
        "kms:Decrypt",
        "kms:GenerateDataKeyWithoutPlaintext",
        "kms:Encrypt",
        "kms:ReEncryptFrom",
        "kms:ReEncryptTo",
        "kms:DescribeKey"
    ],
    "Resource": "*"
}
```

Understanding Source Context

Source context provides information on the source caller attempting to make AWS KMS actions against a given key. This prevents confusion or misuse of encrypted data by binding context to the source of the data.

Customers can utilize source context as additional conditions on their key policy such as the following key policy statements:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "Enable this account full access to this key",
            "Effect": "Allow",
            "Principal": {
                "AWS": "arn:aws:iam::111122223333:root"
            },
            "Action": "kms:*",
            "Resource": "*"
        },
        {
            "Sid": "Enable AVP to retrieve this key's metadata",
            "Effect": "Allow",
            "Principal": {
                "Service": "verifiedpermissions.amazonaws.com"
            },
            "Action": "kms:DescribeKey",
            "Resource": "*",
            "Condition": {
                "StringEquals": {
                    "aws:SourceAccount": "111122223333"
                },
                "StringLike": {
                    "aws:SourceArn": "arn:aws:verifiedpermissions::111122223333:policy-
store/*"
                }
            }
        },
        {
            "Sid": "Enable AVP to encrypt/decrypt resources utilizing this key",
            "Effect": "Allow",
            "Principal": {
```

```
                "Service": "verifiedpermissions.amazonaws.com"
            },
            "Action": [
                "kms:Decrypt",
                "kms:ReEncryptTo",
                "kms:ReEncryptFrom",
                "kms:GenerateDataKeyWithoutPlaintext",
                "kms:Encrypt"
            ],
            "Resource": "*",
            "Condition": {
                "StringEquals": {
                    "aws:SourceAccount": "111122223333"
                },
                "StringLike": {
                    "aws:SourceArn": "arn:aws:verifiedpermissions::111122223333:policy-
store/*"
                }
            }
        }
    ]
}
```

This key policy allows Verified Permissions to make AWS KMS calls on your behalf, if the source account is the same as the account that this AWS KMS key lives in. These values should be verifiable when checking AWS CloudTrail audit logs for the CMK key. For more information on global AWS condition keys, see Using `aws:SourceArn` or `aws:SourceAccount` condition keys.

Understanding Encryption Context

Encryption context is a set of key-value pairs that contain additional authenticated data for encryption integrity checks. When you include an encryption context in a request to encrypt data, AWS KMS cryptographically binds the encryption context to the encrypted data. In order to decrypt the data, you must pass the same encryption context.

Amazon Verified Permissions uses the same encryption context in all AWS KMS cryptographic operations and can be verified within AWS CloudTrail logs when Verified Permissions makes AWS KMS calls on your behalf for encryption/decryption processes. By default, Verified Permissions utilizes the following encryption context key-value pairs when encrypting your resources:

```
{
```

```
    "aws:verifiedpermissions:policy-store-arn":
 "arn:aws:verifiedpermissions::111122223333:policy-store/PSt123456789012"
}
```

Amazon Verified Permissions also allows for you to append custom encryption context as part of additional metadata you wish to include during encryption/decryption processes. This means that your key policy can be more fine-grained in granting permissions such as the example below:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "Enable this account full access to this key",
            "Effect": "Allow",
            "Principal": {
                "AWS": "arn:aws:iam::111122223333:root"
            },
            "Action": "kms:*",
            "Resource": "*"
        },
        {
            "Sid": "Enable AVP to retrieve this key's metadata",
            "Effect": "Allow",
            "Principal": {
                "Service": "verifiedpermissions.amazonaws.com"
            },
            "Action": "kms:DescribeKey",
            "Resource": "*"
        },
        {
            "Sid": "Enable AVP to encrypt/decrypt resources utilizing this key",
            "Effect": "Allow",
            "Principal": {
                "Service": "verifiedpermissions.amazonaws.com"
            },
            "Action": [
                "kms:Decrypt",
                "kms:ReEncryptTo",
                "kms:ReEncryptFrom",
                "kms:GenerateDataKeyWithoutPlaintext",
                "kms:Encrypt"
            ],
            "Resource": "*",
```

```
            "Condition": {
                "StringLike": {
                    "kms:EncryptionContext:aws:verifiedpermissions:policy-store-arn":
 "arn:aws:verifiedpermissions::111122223333:policy-store/*",
                    "kms:EncryptionContext:policy_owner": "Tim"
                }
            }
        }
    ]
}
```

This key policy allows Verified Permissions to make AWS KMS calls on your behalf, if the encryption context map contains a key `aws:verifiedpermissions:policy-store-arn`, whose value follows the format of `arn:aws:verifiedpermissions::111122223333:policy-store/*` and also contains a key-value pair `"policy_owner": "Tim"`. See the section called “Creating an Encrypted Policy store” for how to set custom encryption context.

> **Note**
>
> It is recommended for key policies with conditions based on encryption context to be for a subset of the encryption context map, rather than checking for each key-value pair. The service and its dependencies upstream may add additional key-value pairs that are not visible to you, and can affect Verified Permissions' key access if the key policy conditionally allows based on the **exact** look of the encryption context map.

Understanding kms:ViaService

The `kms:ViaService` condition key limits use of an AWS KMS key to requests from specified AWS services. This condition key only applies for Forward access sessions (FAS). For more information on `kms:ViaService`, see kms:ViaService in the *AWS KMS Developer Guide*.

For example, the following key policy statement uses the `kms:ViaService` condition key to allow a customer managed key to be used for the specified actions only when the request comes from Amazon Verified Permissions in the US East (N. Virginia) region on behalf of `BrentRole`.

```
{
    "Sid": "Enable AVP to encrypt/decrypt resources using credentials of BrentRole",
    "Effect": "Allow",
    "Principal": {
        "AWS": "arn:aws:iam::111122223333:role/BrentRole"
```

```
    },
    "Action": [
        "kms:Decrypt",
        "kms:GenerateDataKeyWithoutPlaintext",
        "kms:Encrypt",
        "kms:ReEncryptFrom",
        "kms:ReEncryptTo",
        "kms:DescribeKey"
    ],
    "Resource": "*",
    "Condition": {
        "StringEquals": {
            "kms:ViaService": [
                "verifiedpermissions.us-east-1.amazonaws.com"
            ]
        }
    }
}
```

This is necessary for Verified Permissions to be able to pass your identity, permissions, and session attributes when Verified Permissions makes a request to AWS KMS on your behalf for encryption/decryption. For more information on FAS requests, see Forward Access Sessions in the *IAM User Guide*.

Complete AWS KMS Key Policy

Based on the concepts in the previous sections, this is an example key policy that will allow Amazon Verified Permissions to use a CMK for encryption/decryption:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "Enable this account full access to this key",
            "Effect": "Allow",
            "Principal": {
                "AWS": "arn:aws:iam::111122223333:root"
            },
            "Action": "kms:*",
            "Resource": "*"
        },
        {
            "Sid": "Enable AVP to retrieve this key's metadata",
```

```
            "Effect": "Allow",
            "Principal": {
                "Service": "verifiedpermissions.amazonaws.com"
            },
            "Action": "kms:DescribeKey",
            "Resource": "*",
            "Condition": {
                "StringEquals": {
                    "aws:SourceAccount": "111122223333"
                },
                "StringLike": {
                    "aws:SourceArn": "arn:aws:verifiedpermissions::111122223333:policy-
store/*"
                }
            }
        },
        {
            "Sid": "Enable AVP to encrypt/decrypt resources utilizing this key",
            "Effect": "Allow",
            "Principal": {
                "Service": "verifiedpermissions.amazonaws.com"
            },
            "Action": [
                "kms:Decrypt",
                "kms:ReEncryptTo",
                "kms:ReEncryptFrom",
                "kms:Encrypt",
                "kms:GenerateDataKeyWithoutPlaintext"
            ],
            "Resource": "*",
            "Condition": {
                "StringLike": {
                    "kms:EncryptionContext:aws:verifiedpermissions:policy-store-arn":
 "arn:aws:verifiedpermissions::111122223333:policy-store/*",
                    "kms:EncryptionContext:policy_owner": "Tim",
                    "aws:SourceArn": "arn:aws:verifiedpermissions::111122223333:policy-
store/*"
                },
                "StringEquals": {
                    "aws:SourceAccount": "111122223333"
                }
            }
        },
        {
```

```
            "Sid": "Enable AVP to encrypt/decrypt resources using credentials of
 BrentRole",
            "Effect": "Allow",
            "Principal": {
                "AWS": "arn:aws:iam::111122223333:role/BrentRole"
            },
            "Action": [
                "kms:Decrypt",
                "kms:GenerateDataKeyWithoutPlaintext",
                "kms:Encrypt",
                "kms:ReEncryptFrom",
                "kms:ReEncryptTo",
                "kms:DescribeKey"
            ],
            "Resource": "*",
            "Condition": {
                "StringEquals": {
                    "kms:ViaService": [
                        "verifiedpermissions.us-east-1.amazonaws.com"
                    ]
                },
                "StringLike": {
                    "kms:EncryptionContext:aws:verifiedpermissions:policy-store-arn":
 "arn:aws:verifiedpermissions::111122223333:policy-store/*",
                    "kms:EncryptionContext:policy_owner": "Tim"
                }
            }
        }
    ]
}
```

Warning

Exercise caution when modifying AWS KMS key policies for keys already in use by Amazon Verified Permissions. While Verified Permissions validates encryption and decryption permissions when you initially configure an AWS KMS key during top-level resource creation, it cannot verify subsequent policy changes on demand. Inadvertently removing necessary permissions could disrupt your authorization decisions and regular Verified Permissions service flows. For guidance troubleshooting common errors related to customer managed keys in Amazon Verified Permissions, refer to the section called "Troubleshoot Customer Managed Keys in Amazon Verified Permissions".

Necessary IAM Policies for Encrypted Resources

Customers that call Verified Permissions via an IAM role within their account will need to ensure that the corresponding IAM policy has proper permissions to utilize the customer managed key for encryption and decryption of resources.

For creating policy stores that are encrypted by a customer managed key, the following IAM policy illustrates the bare-minimum necessary AWS KMS and Verified Permissions actions to do so:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Action": "verifiedpermissions:CreatePolicyStore",
            "Resource": "*",
            "Effect": "Allow"
        },
        {
            "Action": [
                "kms:Decrypt",
                "kms:Encrypt",
                "kms:ReEncryptTo",
                "kms:ReEncryptFrom",
                "kms:DescribeKey",
                "kms:GenerateDataKeyWithoutPlaintext"
            ],
            "Resource": "*",
            "Effect": "Allow"
        }
    ]
}
```

Note

For retrieving (Get* and List* operations) and deleting policy stores that are encrypted by a customer managed key, no additional permissions are needed.

For updating a policy store encrypted by a customer managed key, retrieving (Get* and List* operations), updating, and deleting child resources of a policy store encrypted by a customer

managed key, the following IAM policy illustrates the bare-minimum necessary AWS KMS and Verified Permissions actions do to so:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Action": "verifiedpermissions:*",
            "Resource": "*",
            "Effect": "Allow"
        },
        {
            "Action": [
                "kms:Decrypt"
            ],
            "Resource": "*",
            "Effect": "Allow"
        }
    ]
}
```

As a single IAM policy, customers can simply add the following to their IAM role policy:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Action": "verifiedpermissions:*",
            "Resource": "*",
            "Effect": "Allow"
        },
        {
            "Action": [
                "kms:Decrypt",
                "kms:Encrypt",
                "kms:ReEncryptTo",
                "kms:ReEncryptFrom",
                "kms:DescribeKey",
                "kms:GenerateDataKeyWithoutPlaintext"
            ],
```

```
            "Resource": "*",
            "Effect": "Allow"
        }
    ]
}
```

Managing Encrypted Policy stores

Policy stores are the entry-level container that will contain all related policy resources. For more information about policy stores and the hierarchy of child resources, see Amazon Verified Permissions policy stores in the *Amazon Verified Permissions User Guide*.

When you create a policy store in Verified Permissions, you can enable encryption at rest using AWS KMS keys. This ensures that:

- All read, update, and delete operations on policy stores, and their child resources, will utilize the provided customer managed key for decryption processes
- Any authorization decision calls (i.e. IsAuthorized, BatchIsAuthorized, IsAuthorizedWithToken, etc.) will use the provided customer managed key for decryption processes

Creating an Encrypted Policy store

Before creating an encrypted policy store, ensure that the customer managed key you are using has the proper key policy statements set for Amazon Verified Permissions to utilize the key for encryption/decryption. See the section called “Authorizing use of your AWS KMS key for Amazon Verified Permissions” for what permissions are necessary.

Using AWS CLI:

```
aws verifiedpermissions create-policy-store --region us-east-1 --encryption-settings
 file://encrypted.json --validation-settings "{\"mode\": \"OFF\"}"
```

Where `encrypted.json` looks like:

```
{
    "kmsEncryptionSettings": {
        "key": "arn:aws:kms:us-east-1:111122223333:key/12345678-90ab-cdef-ghij-
klmnopqrstuv",
        "encryptionContext": {
```

```
            "<ENCRYPTION_CONTEXT_KEY_1>": "<ENCRYPTION_CONTEXT_VALUE_1>",
            "<ENCRYPTION_CONTEXT_KEY_2>": "<ENCRYPTION_CONTEXT_VALUE_2>",
            ...
        }
    }
}
```

Making sure to replace `key` with your customer managed key ARN and replacing `<ENCRYPTION_CONTEXT_KEY>` and `<ENCRYPTION_CONTEXT_VALUE>` pairs with the desired `encryptionContext` key-value pairs. `encryptionContext` can be omitted completely if no key-value pair additions are desired.

> **Important**
>
> Do not include the key-value pair `aws:verifiedpermissions:policy-store-arn` in your custom encryption context. This is automatically added and will result in validation errors if it is part of your passed custom encryption context key-value pairs.

For more information of the available APIs of child resources of a policy store, see Actions in the *Amazon Verified Permissions API Reference Guide*.

> **Note**
>
> If the AWS KMS customer managed key in use by your Amazon Verified Permissions resources is deleted, disabled, or inaccessible due to an incorrect AWS KMS key policy, decryption of resources will fail, and thus resulting in stale authorization decisions. The loss of access can be temporary (a key policy can be corrected) or permanent (a deleted key cannot be restored) depending on the circumstances. We recommend you restrict access to critical operations, such as deleting or disabling the AWS KMS key. Also, we recommend that your organization set up AWS break-glass access procedures to ensure your privileged users can access AWS in the unlikely event that Amazon Verified Permissions is inaccessible.

Monitoring Amazon Verified Permissions Interaction with AWS KMS

You can monitor Amazon Verified Permissions' use of your customer managed key through AWS CloudTrail. Each request to AWS KMS via Verified Permissions includes the encryption context and the key ARN being utilized (your customer managed key) in the request parameters:

Example AWS CloudTrail log entry for `GenerateDataKeyWithoutPlaintext`:

```
{
    "eventVersion": "1.11",
    "userIdentity": {
        "type": "AWSService",
        "invokedBy": "verifiedpermissions.amazonaws.com"
    },
    "eventTime": "2025-09-28T16:51:04Z",
    "eventSource": "kms.amazonaws.com",
    "eventName": "GenerateDataKeyWithoutPlaintext",
    "awsRegion": "us-east-1",
    "sourceIPAddress": "verifiedpermissions.amazonaws.com",
    "userAgent": "verifiedpermissions.amazonaws.com",
    "requestParameters": {
        "keyId": "arn:aws:kms:us-east-1:111122223333:key/abcdefgh-0123-ijkl-4567-
mnopqrstuvwx",
        "encryptionContext": {
            "aws:verifiedpermissions:policy-store-arn":
 "arn:aws:verifiedpermissions::111122223333:policy-store/PSt123456789012",
            "policy_store_editor": "Janus"
        },
        ...
    },
    ...
}
```

Example AWS CloudTrail log entry for `Decrypt`:

```
{
    "eventVersion": "1.11",
    "userIdentity": {
        "type": "AWSService",
        "invokedBy": "verifiedpermissions.amazonaws.com"
    },
    "eventTime": "2025-09-28T16:53:21Z",
    "eventSource": "kms.amazonaws.com",
    "eventName": "Decrypt",
    "awsRegion": "us-east-1",
    "sourceIPAddress": "verifiedpermissions.amazonaws.com",
    "userAgent": "verifiedpermissions.amazonaws.com",
    "requestParameters": {
        "encryptionAlgorithm": "SYMMETRIC_DEFAULT",
```

```
        "keyId": "arn:aws:kms:us-east-1:111122223333:key/abcdefgh-0123-ijkl-4567-
mnopqrstuvwx",
        "encryptionContext": {
            "aws:verifiedpermissions:policy-store-arn":
 "arn:aws:verifiedpermissions::111122223333:policy-store/PSt123456789012",
            "policy_store_owner": "Elias"
        }
    },
    ...
}
```

Example AWS CloudTrail log entry for ReEncrypt:

```
{
    "eventVersion": "1.11",
    "userIdentity": {
        "type": "AWSService",
        "invokedBy": "verifiedpermissions.amazonaws.com"
    },
    "eventTime": "2025-09-28T16:51:04Z",
    "eventSource": "kms.amazonaws.com",
    "eventName": "ReEncrypt",
    "awsRegion": "us-east-1",
    "sourceIPAddress": "verifiedpermissions.amazonaws.com",
    "userAgent": "verifiedpermissions.amazonaws.com",
    "requestParameters": {
        "sourceKeyId": "arn:aws:kms:us-east-1:111122223333:key/abcdefgh-0123-ijkl-4567-
mnopqrstuvwx",
        "destinationEncryptionContext": {
            "aws:verifiedpermissions:policy-store-arn":
 "arn:aws:verifiedpermissions::111122223333:policy-store/PSt123456789012"
        },
        "sourceEncryptionAlgorithm": "SYMMETRIC_DEFAULT",
        "destinationKeyId": "arn:aws:kms:us-east-1:111122223333:key/abcdefgh-0123-
ijkl-4567-mnopqrstuvwx",
        "sourceEncryptionContext": {
            "aws:verifiedpermissions:policy_store_arn":
 "arn:aws:verifiedpermissions::111122223333:policy-store/PSt123456789012"
        },
        "destinationEncryptionAlgorithm": "SYMMETRIC_DEFAULT",
        ...
    },
    ...
```

```
}
```

Notice that the log entries include `invokedBy` referencing Amazon Verified Permissions' principal, and `encryptionContext/sourceEncryptionContext/destinationEncryptionContext` being included in the `requestParameters` map.

Example AWS CloudTrail log entry for `DescribeKey`:

```
{
    "eventVersion": "1.11",
    "userIdentity": {
        "type": "AWSService",
        "invokedBy": "verifiedpermissions.amazonaws.com"
    },
    "eventTime": "2025-09-28T16:51:02Z",
    "eventSource": "kms.amazonaws.com",
    "eventName": "DescribeKey",
    "awsRegion": "us-east-1",
    "sourceIPAddress": "verifiedpermissions.amazonaws.com",
    "userAgent": "verifiedpermissions.amazonaws.com",
    "requestParameters": {
        "keyId": "arn:aws:kms:us-east-1:111122223333:key/abcdefgh-0123-ijkl-4567-
mnopqrstuvwx"
    },
    ...
}
```

Notice that the log entry includes `invokedBy` referencing Amazon Verified Permissions' principal.

For more information on AWS CloudTrail log entries, see Understanding AWS CloudTrail events in the *AWS CloudTrail User Guide*.

Limitations

This topic describes the current limitations of Verified Permissions and utilizing customer managed keys for encryption of resources.

- You cannot disable encryption for a policy store once enabled
- After you create a policy store without encryption, you cannot update the policy store to be encrypted by a customer managed key

- After you revoke Verified Permissions access to a customer managed key for an existing encrypted policy store, there is a potential for stale authorization decisions
- After you create a policy store with a customer managed key, you cannot modify custom encryption context values; they are static values set during encrypted policy store creation

Troubleshoot Customer Managed Keys in Amazon Verified Permissions

This topic describes common customer managed key related errors you might encounter when using Amazon Verified Permissions and provides troubleshooting steps to resolve them.

Access Denied: AWS KMS Permission Issue

Error: "Service or caller is not authorized to use the provided AWS KMS key, because the resource does not exist in this Region, no resource-based policies allow access, or a resource-based policy explicitly denies access"

This could either mean that the service or the caller lacks the required `kms:*` action(s) permissions in their IAM policy/AWS KMS key policy or that the key being referenced does not exist or no longer exists.

Troubleshooting with AWS CloudTrail:

- Look for `kms.amazonaws.com` events in AWS CloudTrail
- Search for event name of the AWS KMS operation that was identified to not be allowed (i.e. `Decrypt`, `ReEncrypt`, `GenerateDataKeyWithoutPlaintext`, `DescribeKey`, etc.)
- Review the `errorCode` and `errorMessage` fields
- Check `userIdentity` to confirm which principal attempted the operation

To resolve this issue, grant the user or IAM principal the proper AWS KMS operation access permissions in their IAM policy and AWS KMS key policy. For more information, see the section called “Complete AWS KMS Key Policy”.

Validation Exception: AWS KMS Key Configuration

Error: "Configured AWS KMS key does not have a valid configuration"

This means that the key being referenced cannot be used by the service for customer managed key encryption due to its current configuration. Reasons might include the key being disabled, the key has an unsupported EncryptionAlgorithm, or the key has an unsupported KeyUsage type.

Throttling Exception: AWS KMS Rate Limits

Error: "You have exceeded the rate at which you may call AWS KMS"

This error means that you have exceeded the AWS KMS limit for cryptographic operations for your key: https://docs.aws.amazon.com/kms/latest/developerguide/requests-per-second.html.

Related Information

- Managing Verified Permissions Policy stores
- AWS KMS Best Practices
- AWS KMS Encryption Context
- AWS CloudTrail Integration
- AWS CloudTrail Log Entry Examples

Identity and access management for Amazon Verified Permissions

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 Verified Permissions resources. IAM is an AWS service that you can use with no additional charge.

Topics

Audience

How you use AWS Identity and Access Management (IAM) differs based on your role:

- **Service user** - request permissions from your administrator if you cannot access features (see Troubleshooting Amazon Verified Permissions identity and access)
- **Service administrator** - determine user access and submit permission requests (see How Amazon Verified Permissions works with IAM)
- **IAM administrator** - write policies to manage access (see Identity-based policy examples for Amazon Verified Permissions)

Authenticating with identities

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*.

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*.

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 within your AWS account that has specific permissions. It is similar to an IAM user, but is not associated with a specific person. You can temporarily assume an IAM role in the AWS Management Console by switching roles. You can assume a role by calling an AWS CLI or AWS API operation or by using a custom URL. For more information about methods for using roles, see Using IAM roles in the *IAM User Guide*.

IAM roles with temporary credentials are useful in the following situations:

- **Federated user access** – To assign permissions to a federated identity, you create a role and define permissions for the role. When a federated identity authenticates, the identity is associated with the role and is granted the permissions that are defined by the role. For information about roles for federation, see Create a role for a third-party identity provider (federation) in the *IAM User Guide*. If you use IAM Identity Center, you configure a permission set. To control what your identities can access after they authenticate, IAM Identity Center correlates the permission set to a role in IAM. For information about permissions sets, see Permission sets in the *AWS IAM Identity Center User Guide*.
- **Temporary IAM user permissions** – An IAM user or role can assume an IAM role to temporarily take on different permissions for a specific task.
- **Cross-account access** – You can use an IAM role to allow someone (a trusted principal) in a different account to access resources in your account. Roles are the primary way to grant cross-account access. However, with some AWS services, you can attach a policy directly to a resource (instead of using a role as a proxy). To learn the difference between roles and resource-based policies for cross-account access, see How IAM roles differ from resource-based policies in the *IAM User Guide*.

- **Applications running on Amazon EC2** – You can use an IAM role to manage temporary credentials for applications that are running on an EC2 instance and making AWS CLI or AWS API requests. This is preferable to storing access keys within the EC2 instance. To assign an AWS role to an EC2 instance and make it available to all of its applications, you create an instance profile that is attached to the instance. An instance profile contains the role and enables programs that are running on the EC2 instance to get temporary credentials. For more information, see Use an IAM role to grant permissions to applications running on Amazon EC2 instances in the *IAM User Guide*.

To learn whether to use IAM roles or IAM users, see When to create an IAM role (instead of a user) in the *IAM User 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 attach to an identity (user, group, or role). These policies control what actions identities 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 *inline policies* (embedded directly into a single identity) or *managed policies* (standalone policies attached to multiple identities). To learn how to choose between managed and inline policies, see Choose between managed policies and inline policies in the *IAM User Guide*.

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*.

How Amazon Verified Permissions works with IAM

Before you use IAM to manage access to Verified Permissions, learn what IAM features are available to use with Verified Permissions.

IAM features you can use with Amazon Verified Permissions

IAM feature	Verified Permissions support
Identity-based policies	Yes
Resource-based policies	No
Policy actions	Yes
Policy resources	Yes
Policy condition keys	No
ACLs	No
ABAC (tags in policies)	Yes
Temporary credentials	Yes
Principal permissions	Yes
Service roles	No
Service-linked roles	No

To get a high-level view of how Verified Permissions and other AWS services work with most IAM features, see AWS services that work with IAM in the *IAM User Guide*.

Identity-based policies for Verified Permissions

Supports identity-based policies	Yes

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles 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*.

With IAM identity-based policies, you can specify allowed or denied actions and resources as well as the conditions under which actions are allowed or denied. To learn about all of the elements that you can use in a JSON policy, see IAM JSON policy elements reference in the *IAM User Guide*.

Identity-based policy examples for Verified Permissions

To view examples of Verified Permissions identity-based policies, see Identity-based policy examples for Amazon Verified Permissions.

Resource-based policies within Verified Permissions

Supports resource-based policies	No

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are 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. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must specify a principal in a resource-based policy. Principals can include accounts, users, roles, federated users, or AWS services.

To enable cross-account access, you can specify an entire account or IAM entities in another account as the principal in a resource-based policy. For more information, see Cross account resource access in IAM in the *IAM User Guide*.

Policy actions for Verified Permissions

Supports policy actions	Yes

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Action` element of a JSON policy describes the actions that you can use to allow or deny access in a policy. Include actions in a policy to grant permissions to perform the associated operation.

To see a list of Verified Permissions actions, see Actions defined by Amazon Verified Permissions in the *Service Authorization Reference*.

Policy actions in Verified Permissions use the following prefix before the action:

```
verifiedpermissions
```

To specify multiple actions in a single statement, separate them with commas.

```
"Action": [
      "verifiedpermissions:action1",
      "verifiedpermissions:action2"
         ]
```

You can specify multiple actions using wildcards (*). For example, to specify all actions that begin with the word `Get`, include the following action:

```
"Action": "verifiedpermissions:Get*"
```

To view examples of Verified Permissions identity-based policies, see Identity-based policy examples for Amazon Verified Permissions.

Policy resources for Verified Permissions

Supports policy resources	Yes

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Resource` JSON policy element specifies the object or objects to which the action applies. As a best practice, specify a resource using its Amazon Resource Name (ARN). For actions that don't support resource-level permissions, use a wildcard (*) to indicate that the statement applies to all resources.

```
"Resource": "*"
```

To see a list of Verified Permissions resource types and their ARNs, see Resource types defined by Amazon Verified Permissions in the *Service Authorization Reference*. To learn with which actions you can specify the ARN of each resource, see Actions defined by Amazon Verified Permissions.

Policy condition keys for Verified Permissions

Supports service-specific policy condition keys	No

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Condition` element specifies when statements execute based on defined criteria. You can create conditional expressions that use condition operators, such as equals or less than, to match the condition in the policy with values in the request. To see all AWS global condition keys, see AWS global condition context keys in the *IAM User Guide*.

ACLs in Verified Permissions

Supports ACLs	No

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.

ABAC with Verified Permissions

Supports ABAC (tags in policies)	Yes

Attribute-based access control (ABAC) is an authorization strategy that defines permissions based on attributes called tags. You can attach tags to IAM entities and AWS resources, then design ABAC policies to allow operations when the principal's tag matches the tag on the resource.

To control access based on tags, you provide tag information in the condition element of a policy using the `aws:ResourceTag/key-name`, `aws:RequestTag/key-name`, or `aws:TagKeys` condition keys.

If a service supports all three condition keys for every resource type, then the value is **Yes** for the service. If a service supports all three condition keys for only some resource types, then the value is **Partial**.

For more information about ABAC, see Define permissions with ABAC authorization in the *IAM User Guide*. To view a tutorial with steps for setting up ABAC, see Use attribute-based access control (ABAC) in the *IAM User Guide*.

Using temporary credentials with Verified Permissions

Supports temporary credentials	Yes

Temporary credentials provide short-term access to AWS resources and are automatically created when you use federation or switch roles. AWS recommends that you dynamically generate temporary credentials instead of using long-term access keys. For more information, see Temporary security credentials in IAM and AWS services that work with IAM in the *IAM User Guide*.

Cross-service principal permissions for Verified Permissions

Supports principal permissions	Yes

Forward access sessions (FAS) use the permissions of the principal calling an AWS service, combined with the requesting AWS service to make requests to downstream services. For policy details when making FAS requests, see Forward access sessions.

Service roles for Verified Permissions

Supports service roles	No

A service role is an IAM role that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For more information, see Create a role to delegate permissions to an AWS service in the *IAM User Guide*.

Service-linked roles for Verified Permissions

Supports service-linked roles	No

A service-linked role is a type of service role that is linked to an AWS service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your AWS account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.

For details about creating or managing service-linked roles, see AWS services that work with IAM. Find a service in the table that includes a `Yes` in the **Service-linked role** column. Choose the **Yes** link to view the service-linked role documentation for that service.

IAM policies for Verified Permissions

Verified Permissions manages the permissions of users within your application. In order for your application to call the Verified Permissions APIs or for AWS Management Console users to be allowed to manage Cedar policies in a Verified Permissions policy store, you must add the necessary IAM permissions.

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see Creating IAM policies in the IAM User Guide.

With IAM identity-based policies, you can specify allowed or denied actions and resources as well as the conditions under which actions are allowed or denied (listed below). You can't specify the principal in an identity-based policy because it applies to the user or role to which it is attached. To

learn about all of the elements that you can use in a JSON policy, see IAM JSON policy elements reference in the IAM User Guide.

Action	Description
CreateIdentitySource	Action to create a new identity source.
CreatePolicy	Action to create a Cedar policy in a policy store. You can create either a static policy or a policy linked to a policy template.
CreatePolicyStore	Action to create a new policy store.
CreatePolicyTemplate	Action to create a new policy template.
DeleteIdentitySource	Action to delete an identity source.
DeletePolicy	Action to delete a policy from a policy store.
DeletePolicyStore	Action to delete a policy store.
DeletePolicyTemplate	Action to delete a policy template.
GetIdentitySource	Action to get an identity source.
GetPolicy	Action to retrieve information about a specified policy.
GetPolicyStore	Action to retrieve information about a specified policy store.
GetPolicyTemplate	Action to get a policy template.
GetSchema	Action to get a schema.
IsAuthorized	Action to get an authorization response based on the parameters described in the authorization request.
IsAuthorizedWithToken	Action to get an authorization response based on the parameters described in the authoriza

Action	Description
	tion request where the principal comes from an identity token.
ListIdentitySources	Action to list all the identity sources in the AWS account.
ListPolicies	Action to list all policies in a policy store.
ListPolicyStores	Action to list all policy stores in the AWS account.
ListPolicyTemplates	Action to list all policy templates in the AWS account.
ListTagsForResource	Action to list all the tags for a resource.
PutSchema	Action to add a schema to a policy store.
TagResource	Action to add a tag to a resource.
UpdateIdentitySource	Action to update an identity source.
UpdatePolicy	Action to update a policy in a policy store.
UpdatePolicyStore	Action to update a policy store.
UpdatePolicyTemplate	Action to update a policy template.
UntagResource	Action to remove a tag from a resource.

Example IAM policy for permission to the CreatePolicy action:

JSON

```
{
"Version":"2012-10-17",
    "Statement": [
        {
```

```
            "Effect": "Allow",
            "Action": [
                "verifiedpermissions:CreatePolicy"
            ],
            "Resource": "*"
        }
    ]
}
```

Identity-based policy examples for Amazon Verified Permissions

By default, users and roles don't have permission to create or modify Verified Permissions resources. They also can't perform tasks by using the AWS Management Console, AWS Command Line Interface (AWS CLI), or AWS API. An IAM administrator must create IAM policies that grant users and roles permission to perform actions on the resources that they need. The administrator must then attach those policies for users that require them.

To learn how to create an IAM identity-based policy by using these example JSON policy documents, see Creating IAM policies in the *IAM User Guide*.

For details about actions and resource types defined by Verified Permissions, including the format of the ARNs for each of the resource types, see see Actions, resources, and condition keys for Amazon Verified Permissions in the *Service Authorization Reference*.

Topics

- Policy best practices
- Using the Verified Permissions console
- Allow users to view their own permissions

Policy best practices

Identity-based policies determine whether someone can create, access, or delete Verified Permissions resources in your account. These actions can incur costs for your AWS account. When you create or edit identity-based policies, follow these guidelines and recommendations:

- **Get started with AWS managed policies and move toward least-privilege permissions** – To get started granting permissions to your users and workloads, use the *AWS managed policies* that grant permissions for many common use cases. They are available in your AWS account. We

recommend that you reduce permissions further by defining AWS customer managed policies that are specific to your use cases. For more information, see AWS managed policies or AWS managed policies for job functions in the *IAM User Guide*.

- **Apply least-privilege permissions** – When you set permissions with IAM policies, grant only the permissions required to perform a task. You do this by defining the actions that can be taken on specific resources under specific conditions, also known as *least-privilege permissions*. For more information about using IAM to apply permissions, see Policies and permissions in IAM in the *IAM User Guide*.
- **Use conditions in IAM policies to further restrict access** – You can add a condition to your policies to limit access to actions and resources. For example, you can write a policy condition to specify that all requests must be sent using SSL. You can also use conditions to grant access to service actions if they are used through a specific AWS service, such as CloudFormation. For more information, see IAM JSON policy elements: Condition in the *IAM User Guide*.
- **Use IAM Access Analyzer to validate your IAM policies to ensure secure and functional permissions** – IAM Access Analyzer validates new and existing policies so that the policies adhere to the IAM policy language (JSON) and IAM best practices. IAM Access Analyzer provides more than 100 policy checks and actionable recommendations to help you author secure and functional policies. For more information, see Validate policies with IAM Access Analyzer in the *IAM User Guide*.
- **Require multi-factor authentication (MFA)** – If you have a scenario that requires IAM users or a root user in your AWS account, turn on MFA for additional security. To require MFA when API operations are called, add MFA conditions to your policies. For more information, see Secure API access with MFA in the *IAM User Guide*.

For more information about best practices in IAM, see Security best practices in IAM in the *IAM User Guide*.

Using the Verified Permissions console

To access the Amazon Verified Permissions console, you must have a minimum set of permissions. These permissions must allow you to list and view details about the Verified Permissions resources in your AWS account. If you create an identity-based policy that is more restrictive than the minimum required permissions, the console won't function as intended for entities (users or roles) with that policy.

You don't need to allow minimum console permissions for users that are making calls only to the AWS CLI or the AWS API. Instead, allow access to only the actions that match the API operation that they're trying to perform.

To ensure that users and roles can still use the Verified Permissions console, also attach the Verified Permissions *ConsoleAccess* or *ReadOnly* AWS managed policy to the entities. For more information, see Adding permissions to a user in the *IAM User Guide*.

Allow users to view their own permissions

This example shows how you might create a policy that allows IAM users to view the inline and managed policies that are attached to their user identity. This policy includes permissions to complete this action on the console or programmatically using the AWS CLI or AWS API.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "ViewOwnUserInfo",
            "Effect": "Allow",
            "Action": [
                "iam:GetUserPolicy",
                "iam:ListGroupsForUser",
                "iam:ListAttachedUserPolicies",
                "iam:ListUserPolicies",
                "iam:GetUser"
            ],
            "Resource": ["arn:aws:iam::*:user/${aws:username}"]
        },
        {
            "Sid": "NavigateInConsole",
            "Effect": "Allow",
            "Action": [
                "iam:GetGroupPolicy",
                "iam:GetPolicyVersion",
                "iam:GetPolicy",
                "iam:ListAttachedGroupPolicies",
                "iam:ListGroupPolicies",
                "iam:ListPolicyVersions",
                "iam:ListPolicies",
                "iam:ListUsers"
            ],
            "Resource": "*"
```

```
        }
    ]
}
```

AWS managed policies for Amazon Verified Permissions

To add permissions to users, groups, and roles, it is easier to use AWS managed policies than to write policies yourself. It takes time and expertise to create IAM customer managed policies that provide your team with only the permissions they need. To get started quickly, you can use our AWS managed policies. These policies cover common use cases and are available in your AWS account. For more information about AWS managed policies, see AWS managed policies in the *IAM User Guide*.

AWS services maintain and update AWS managed policies. You can't change the permissions in AWS managed policies. Services occasionally add additional permissions to an AWS managed policy to support new features. This type of update affects all identities (users, groups, and roles) where the policy is attached. Services are most likely to update an AWS managed policy when a new feature is launched or when new operations become available. Services do not remove permissions from an AWS managed policy, so policy updates won't break your existing permissions.

Additionally, AWS supports managed policies for job functions that span multiple services. For example, the **ReadOnlyAccess** AWS managed policy provides read-only access to all AWS services and resources. When a service launches a new feature, AWS adds read-only permissions for new operations and resources. For a list and descriptions of job function policies, see AWS managed policies for job functions in the *IAM User Guide*.

AWS managed policy: AmazonVerifiedPermissionsFullAccess

The `AmazonVerifiedPermissionsFullAccess` managed policy grants full access to Verified Permissions. To work with Amazon Cognito-based identity sources, you'll need to attach a separate policy, such as the AmazonCognitoReadOnly policy.

JSON

```
{
```

```
  "Version":"2012-10-17",
  "Statement": [
    {
      "Sid": "AccountLevelPermissions",
      "Effect": "Allow",
      "Action": [
        "verifiedpermissions:CreatePolicyStore",
        "verifiedpermissions:ListPolicyStores"
      ],
      "Resource": "*"
    },
    {
      "Sid": "PolicyStoreLevelPermissions",
      "Effect": "Allow",
      "Action": [
        "verifiedpermissions:*"
      ],
      "Resource": [
        "arn:aws:verifiedpermissions::*:policy-store/*"
      ]
    }
  ]
}
```

AWS managed policy: AmazonVerifiedPermissionsReadOnlyAccess

The `AmazonVerifiedPermissionsReadOnlyAccess` managed policy grants read-only access to Verified Permissions.

This policy grants access to all read operations of Amazon Verified Permissions, including the authorization query APIs `IsAuthorized` and `IsAuthorizedWithToken`.

> **Note**
>
> Access to `BatchIsAuthorized` and `BatchIsAuthorizedWithToken` are granted automatically when access is granted to `IsAuthorized` and `IsAuthorizedWithToken`, respectively.

JSON

```
{
  "Version":"2012-10-17",
  "Statement": [
    {
      "Sid": "AccountLevelPermissions",
      "Effect": "Allow",
      "Action": [
        "verifiedpermissions:ListPolicyStores"
      ],
      "Resource": "*"
    },
    {
      "Sid": "PolicyStoreLevelPermissions",
      "Effect": "Allow",
      "Action": [
        "verifiedpermissions:GetIdentitySource",
        "verifiedpermissions:GetPolicy",
        "verifiedpermissions:GetPolicyStore",
        "verifiedpermissions:GetPolicyTemplate",
        "verifiedpermissions:GetSchema",
        "verifiedpermissions:IsAuthorized",
        "verifiedpermissions:IsAuthorizedWithToken",
        "verifiedpermissions:ListIdentitySources",
        "verifiedpermissions:ListPolicies",
        "verifiedpermissions:ListPolicyTemplates"
      ],
      "Resource": [
        "arn:aws:verifiedpermissions::*:policy-store/*"
      ]
    }
  ]
}
```

Verified Permissions updates to AWS managed policies

View details about updates to AWS managed policies for Verified Permissions since this service began tracking these changes. For automatic alerts about changes to this page, subscribe to the RSS feed on the Verified Permissions Document history page.

Change	Description	Date
AmazonVerifiedPermissionsFullAccess – New policy	Verified Permissions added a new policy to allow full access to Verified Permissions.	October 11, 2024
AmazonVerifiedPermissionsReadOnlyAccess – New policy	Verified Permissions added a new policy to allow access to all read operations of Amazon Verified Permissions, including the authorization query APIs `IsAuthorized` and `IsAuthorizedWithToken` .	October 11, 2024
Verified Permissions started tracking changes	Verified Permissions started tracking changes for its AWS managed policies.	October 11, 2024

Troubleshooting Amazon Verified Permissions identity and access

Use the following information to help you diagnose and fix common issues that you might encounter when working with Verified Permissions and IAM.

Topics

- I am not authorized to perform an action in Verified Permissions
- I am not authorized to perform iam:PassRole
- I want to allow people outside of my AWS account to access my Verified Permissions resources

I am not authorized to perform an action in Verified Permissions

If you receive an error that you're not authorized to perform an action, your policies must be updated to allow you to perform the action.

The following example error occurs when the `mateojackson` IAM user tries to use the console to view details about a fictional *`my-example-widget`* resource but doesn't have the fictional `verifiedpermissions:`*`GetWidget`* permissions.

```
User: arn:aws:iam::123456789012:user/mateojackson is not authorized to perform:
 verifiedpermissions:GetWidget on resource: my-example-widget
```

In this case, the policy for the `mateojackson` user must be updated to allow access to the *`my-example-widget`* resource by using the `verifiedpermissions:`*`GetWidget`* action.

If you need help, contact your AWS administrator. Your administrator is the person who provided you with your sign-in credentials.

I am not authorized to perform iam:PassRole

If you receive an error that you're not authorized to perform the `iam:PassRole` action, your policies must be updated to allow you to pass a role to Verified Permissions.

Some AWS services allow you to pass an existing role to that service instead of creating a new service role or service-linked role. To do this, you must have permissions to pass the role to the service.

The following example error occurs when an IAM user named `marymajor` tries to use the console to perform an action in Verified Permissions. However, the action requires the service to have permissions that are granted by a service role. Mary does not have permissions to pass the role to the service.

```
User: arn:aws:iam::123456789012:user/marymajor is not authorized to perform:
 iam:PassRole
```

In this case, Mary's policies must be updated to allow her to perform the `iam:PassRole` action.

If you need help, contact your AWS administrator. Your administrator is the person who provided you with your sign-in credentials.

I want to allow people outside of my AWS account to access my Verified Permissions resources

You can create a role that users in other accounts or people outside of your organization can use to access your resources. You can specify who is trusted to assume the role. For services that support resource-based policies or access control lists (ACLs), you can use those policies to grant people access to your resources.

To learn more, consult the following:

- To learn whether Verified Permissions supports these features, see How Amazon Verified Permissions works with IAM.
- To learn how to provide access to your resources across AWS accounts that you own, see Providing access to an IAM user in another AWS account that you own in the *IAM User Guide*.
- To learn how to provide access to your resources to third-party AWS accounts, see Providing access to AWS accounts owned by third parties in the *IAM User Guide*.
- To learn how to provide access through identity federation, see Providing access to externally authenticated users (identity federation) in the *IAM User Guide*.
- To learn the difference between using roles and resource-based policies for cross-account access, see Cross account resource access in IAM in the *IAM User Guide*.

Compliance validation for Amazon Verified Permissions

Our new AWS sign-up experience is not designed for regulated workloads. If you're using our new AWS sign-up experience, but you want to use AWS for regulated workloads, you can sign up for AWS (advanced) or activate advanced features for your AWS environment.

To learn whether an AWS service is within the scope of specific compliance programs, see AWS services in Scope by Compliance Program and choose the compliance program that you are interested in. 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 services is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. For more information about your compliance responsibility when using AWS services, see AWS Security Documentation.

Resilience in Amazon Verified Permissions

The AWS global infrastructure is built around AWS Regions and Availability Zones. AWS Regions provide multiple physically separated and isolated Availability Zones, which are connected with 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.

When you create a Verified Permissions policy store , it is created within an individual AWS Region, and is automatically replicated across the data centers that make up that Region's Availability Zones. At this time, Verified Permissions doesn't support any cross-region replication.

For more information about AWS Regions and Availability Zones, see AWS Global Infrastructure.

Monitoring Amazon Verified Permissions API calls

Monitoring is an important part of maintaining the reliability, availability, and performance of Amazon Verified Permissions and your other AWS solutions. AWS provides the following tools to monitor Verified Permissions, report when something is wrong, and take automatic actions when appropriate:

- *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.

For more information about monitoring Verified Permissions with CloudTrail, see Logging Amazon Verified Permissions API calls using AWS CloudTrail.

Logging Amazon Verified Permissions API calls using AWS CloudTrail

Amazon Verified Permissions is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service in Verified Permissions. CloudTrail captures all API calls for Verified Permissions as events. The calls captured include calls from the Verified Permissions console and code calls to the Verified Permissions API operations. If you create a trail, you can enable continuous delivery of CloudTrail events to an Amazon S3 bucket, including events for Verified Permissions. If you don't configure a trail, you can still view the most recent management action events in the CloudTrail console in **Event history**, but not events for API calls such as `isAuthorized`. Using the information collected by CloudTrail, you can determine the request that was made to Verified Permissions, 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.

Verified Permissions information in CloudTrail

CloudTrail is enabled on your AWS account when you create the account. When activity occurs in Verified Permissions, 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 Verified Permissions, create a trail. A *trail* enables 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 Verified Permissions actions are logged by CloudTrail and are documented in the Amazon Verified Permissions API Reference Guide. For example, calls to the `CreateIdentitySource`, `DeletePolicy`, and `ListPolicyStores` actions 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.

Data events like IsAuthorized and IsAuthorizedWithToken are not logged by default when you create a trail or event data store. To record CloudTrail data events, you must explicitly add the supported resources or resource types for which you want to collect activity. For more information, see Data events in the *AWS CloudTrail User Guide*.

Understanding Verified Permissions log file entries

A trail is a configuration that enables 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.

For authorization API calls, the response elements, such as the decision, are included under `additionalEventData` rather than `responseElements`.

Topics

> **Note**
>
> Some fields have been redacted from the examples for data privacy.

IsAuthorized

The following example shows a CloudTrail log entry for an `IsAuthorized` API call:

```
{
    "eventVersion": "1.08",
    "userIdentity": {
  "type": "AssumedRole",
  "principalId": "EXAMPLE_PRINCIPAL_ID",
  "arn": "arn:aws:iam::123456789012:role/ExampleRole",
  "accountId": "123456789012",
  "accessKeyId": "AKIAIOSFODNN7EXAMPLE"
    },
    "eventTime": "2023-11-20T22:55:03Z",
    "eventSource": "verifiedpermissions.amazonaws.com",
    "eventName": "IsAuthorized",
    "awsRegion": "us-west-2",
    "sourceIPAddress": "203.0.113.0",
    "userAgent": "aws-cli/2.11.18 Python/3.11.3 Linux/5.4.241-160.348.amzn2int.x86_64
 exe/x86_64.amzn.2 prompt/off command/verifiedpermissions.is-authorized",
    "requestParameters": {
        "principal": {
            "entityType": "PhotoFlash::User",
            "entityId": "alice"
        },
        "action": {
            "actionType": "PhotoFlash::Action",
            "actionId": "ViewPhoto"
        },
        "resource": {
            "entityType": "PhotoFlash::Photo",
            "entityId": "VacationPhoto94.jpg"
        },
        "policyStoreId": "PSEXAMPLEabcdefg111111"
    },
    "responseElements": null,
    "additionalEventData": {
        "decision": "ALLOW"
    },
    "requestID": "346c4b6a-d12f-46b6-bc06-6c857bd3b28e",
    "eventID": "8a4fed32-9605-45dd-a09a-5ebbf0715bbc",
    "readOnly": true,
    "resources": [
        {
    "accountId": "123456789012",
    "type": "AWS::VerifiedPermissions::PolicyStore",
```

```
    "ARN": "arn:aws:verifiedpermissions::123456789012:policy-store/
PSEXAMPLEabcdefg111111"
        }
    ],
    "eventType": "AwsApiCall",
    "managementEvent": false,
    "recipientAccountId": "123456789012",
    "eventCategory": "Data"
}
```

IsAuthorizedWithToken

The following example shows a CloudTrail log entry for an `IsAuthorizedWithToken` API call:

```
{
    "eventVersion": "1.08",
    "userIdentity": {
  "type": "AssumedRole",
  "principalId": "EXAMPLE_PRINCIPAL_ID",
  "arn": "arn:aws:iam::123456789012:role/ExampleRole",
  "accountId": "123456789012",
  "accessKeyId": "AKIAIOSFODNN7EXAMPLE"
    },
    "eventTime": "2023-11-20T22:55:03Z",
    "eventSource": "verifiedpermissions.amazonaws.com",
    "eventName": "IsAuthorizedWithToken",
    "awsRegion": "us-west-2",
    "sourceIPAddress": "203.0.113.0",
    "userAgent": "aws-cli/2.11.18 Python/3.11.3 Linux/5.4.241-160.348.amzn2int.x86_64
 exe/x86_64.amzn.2 prompt/off command/verifiedpermissions.is-authorized-with-token",
    "requestParameters": {
        "policyStoreId": "PSEXAMPLEabcdefg111111",
        "resource": {
            "entityType": "PhotoFlash::Photo",
            "entityId": "VacationPhoto94.jpg"
        },
        "action": {
            "actionType": "PhotoFlash::Action",
            "actionId": "ViewPhoto"
        }
    },
    "responseElements": null,
    "additionalEventData": {
```

```
        "decision": "ALLOW",
        "principal": {
            "entityType": "PhotoFlash::User",
            "entityId": "us-east-1_EXAMPLE|EXAMPLE_SUBJECT_ID"
        }
    },
    "requestID": "346c4b6a-d12f-46b6-bc06-6c857bd3b28e",
    "eventID": "8a4fed32-9605-45dd-a09a-5ebbf0715bbc",
    "readOnly": true,
    "resources": [
        {
    "accountId": "123456789012",
    "type": "AWS::VerifiedPermissions::PolicyStore",
    "ARN": "arn:aws:verifiedpermissions::123456789012:policy-store/
PSEXAMPLEabcdefg111111"
        }
    ],
    "eventType": "AwsApiCall",
    "managementEvent": false,
    "recipientAccountId": "123456789012",
    "eventCategory": "Data"
}
```

BatchIsAuthorized

The following example shows a CloudTrail log entry for a `BatchIsAuthorized` API call:

```
{
    "eventVersion": "1.08",
    "userIdentity": {
  "type": "AssumedRole",
  "principalId": "EXAMPLE_PRINCIPAL_ID",
  "arn": "arn:aws:iam::123456789012:role/ExampleRole",
  "accountId": "123456789012",
  "accessKeyId": "AKIAIOSFODNN7EXAMPLE"
    },
    "eventTime": "2023-11-20T23:02:33Z",
    "eventSource": "verifiedpermissions.amazonaws.com",
    "eventName": "BatchIsAuthorized",
    "awsRegion": "us-west-2",
    "sourceIPAddress": "203.0.113.0",
    "userAgent": "aws-cli/2.11.18 Python/3.11.3 Linux/5.4.241-160.348.amzn2int.x86_64
 exe/x86_64.amzn.2 prompt/off command/verifiedpermissions.is-authorized",
```

```
    "requestParameters": {
        "requests": [
            {
                "principal": {
                    "entityType": "PhotoFlash::User",
                    "entityId": "alice"
                },
                "action": {
                    "actionType": "PhotoFlash::Action",
                    "actionId": "ViewPhoto"
                },
                "resource": {
                    "entityType": "PhotoFlash::Photo",
                    "entityId": "VacationPhoto94.jpg"
                }
            },
            {
                "principal": {
                    "entityType": "PhotoFlash::User",
                    "entityId": "annalisa"
                },
                "action": {
                    "actionType": "PhotoFlash::Action",
                    "actionId": "DeletePhoto"
                },
                "resource": {
                    "entityType": "PhotoFlash::Photo",
                    "entityId": "VacationPhoto94.jpg"
                }
            }
        ],
        "policyStoreId": "PSEXAMPLEabcdefg111111"
    },
    "responseElements": null,
    "additionalEventData": {
        "results": [
            {
                "request": {
                    "principal": {
                        "entityType": "PhotoFlash::User",
                        "entityId": "alice"
                    },
                    "action": {
                        "actionType": "PhotoFlash::Action",
```

```
                        "actionId": "ViewPhoto"
                    },
                    "resource": {
                        "entityType": "PhotoFlash::Photo",
                        "entityId": "VacationPhoto94.jpg"
                    }
                },
                "decision": "ALLOW"
            },
            {
                "request": {
                    "principal": {
                        "entityType": "PhotoFlash::User",
                        "entityId": "annalisa"
                    },
                    "action": {
                        "actionType": "PhotoFlash::Action",
                        "actionId": "DeletePhoto"
                    },
                    "resource": {
                        "entityType": "PhotoFlash::Photo",
                        "entityId": "VacationPhoto94.jpg"
                    }
                },
                "decision": "DENY"
            }
        ]
    },
    "requestID": "a8a5caf3-78bd-4139-924c-7101a8339c3b",
    "eventID": "7d81232f-f3d1-4102-b9c9-15157c70487b",
    "readOnly": true,
    "resources": [
        {
    "accountId": "123456789012",
    "type": "AWS::VerifiedPermissions::PolicyStore",
    "ARN": "arn:aws:verifiedpermissions::123456789012:policy-store/
PSEXAMPLEabcdefg111111"
        }
    ],
    "eventType": "AwsApiCall",
    "managementEvent": false,
    "recipientAccountId": "123456789012",
    "eventCategory": "Data"
```

```
}
```

BatchIsAuthorizedWithToken

The following example shows a CloudTrail log entry for a `BatchIsAuthorizedWithToken` API call:

```
{
    "eventVersion": "1.08",
    "userIdentity": {
  "type": "AssumedRole",
  "principalId": "EXAMPLE_PRINCIPAL_ID",
  "arn": "arn:aws:iam::123456789012:role/ExampleRole",
  "accountId": "123456789012",
  "accessKeyId": "AKIAIOSFODNN7EXAMPLE"
    },
    "eventTime": "2023-11-20T23:02:33Z",
    "eventSource": "verifiedpermissions.amazonaws.com",
    "eventName": "BatchIsAuthorizedWithToken",
    "awsRegion": "us-west-2",
    "sourceIPAddress": "203.0.113.0",
    "userAgent": "aws-cli/2.11.18 Python/3.11.3 Linux/5.4.241-160.348.amzn2int.x86_64
 exe/x86_64.amzn.2 prompt/off command/verifiedpermissions.batch-is-authorized-with-
token",
    "requestParameters": {
        "policyStoreId": "PSEXAMPLEabcdefg111111",
        "requests": [
            {
                "action": {
                    "actionType": "PhotoFlash::Action",
                    "actionId": "ViewPhoto"
                },
                "resource": {
                    "entityType": "PhotoFlash::Photo",
                    "entityId": "VacationPhoto94.jpg"
                }
            },
            {
                "action": {
                    "actionType": "PhotoFlash::Action",
                    "actionId": "DeletePhoto"
                },
                "resource": {
```

```
                    "entityType": "PhotoFlash::Photo",
                    "entityId": "VacationPhoto94.jpg"
                }
            }
        ]
    },
    "responseElements": null,
    "additionalEventData": {
        "principal": {
            "entityType": "PhotoFlash::User",
            "entityId": "us-east-1_EXAMPLE|EXAMPLE_SUBJECT_ID"
        },
        "results": [
            {
                "request": {
                    "action": {
                        "actionType": "PhotoFlash::Action",
                        "actionId": "ViewPhoto"
                    },
                    "resource": {
                        "entityType": "PhotoFlash::Photo",
                        "entityId": "VacationPhoto94.jpg"
                    }
                },
                "decision": "ALLOW"
            },
            {
                "request": {
                    "action": {
                        "actionType": "PhotoFlash::Action",
                        "actionId": "DeletePhoto"
                    },
                    "resource": {
                        "entityType": "PhotoFlash::Photo",
                        "entityId": "VacationPhoto94.jpg"
                    }
                },
                "decision": "DENY"
            }
        ]
    },
    "requestID": "c67aa1d9-9fd7-45ad-9a23-e8f9f6dd66c8",
    "eventID": "5ed90624-a4f7-4ef6-a07a-0e6486d30fbe",
    "readOnly": true,
```

```
    "resources": [
        {
    "accountId": "123456789012",
    "type": "AWS::VerifiedPermissions::PolicyStore",
    "ARN": "arn:aws:verifiedpermissions::123456789012:policy-store/
PSEXAMPLEabcdefg111111"
        }
    ],
    "eventType": "AwsApiCall",
    "managementEvent": false,
    "recipientAccountId": "123456789012",
    "eventCategory": "Data"
}
```

CreatePolicyStore

The following example shows a CloudTrail log entry for a `CreatePolicyStore` API call:

```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "EXAMPLE_PRINCIPAL_ID",
    "arn": "arn:aws:iam::123456789012:role/ExampleRole",
    "accountId": "123456789012",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE"
  },
  "eventTime": "2023-05-22T07:43:33Z",
  "eventSource": "verifiedpermissions.amazonaws.com",
  "eventName": "CreatePolicyStore",
  "awsRegion": "us-west-2",
  "sourceIPAddress": "203.0.113.0",
  "userAgent": "aws-sdk-rust/0.55.2 os/linux lang/rust/1.69.0",
  "requestParameters": {
    "clientToken": "a1b2c3d4-e5f6-a1b2-c3d4-TOKEN1111111",
    "validationSettings": {
      "mode": "OFF"
    }
  },
  "responseElements": {
    "policyStoreId": "PSEXAMPLEabcdefg111111",
    "arn": "arn:aws:verifiedpermissions::123456789012:policy-store/
PSEXAMPLEabcdefg111111",
```

```
    "createdDate": "2023-05-22T07:43:33.962794Z",
    "lastUpdatedDate": "2023-05-22T07:43:33.962794Z"
  },
  "requestID": "1dd9360e-e2dc-4554-ab65-b46d2cf45c29",
  "eventID": "b6edaeee-3584-4b4e-a48e-311de46d7532",
  "readOnly": false,
  "eventType": "AwsApiCall",
  "managementEvent": true,
  "recipientAccountId": "123456789012",
  "eventCategory": "Management"
}
```

ListPolicyStores

The following example shows a CloudTrail log entry for a `ListPolicyStores` API call:

```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "EXAMPLE_PRINCIPAL_ID",
    "arn": "arn:aws:iam::123456789012:role/ExampleRole",
    "accountId": "123456789012",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE"
  },
  "eventTime": "2023-05-22T07:43:33Z",
  "eventSource": "verifiedpermissions.amazonaws.com",
  "eventName": "ListPolicyStores",
  "awsRegion": "us-west-2",
  "sourceIPAddress": "203.0.113.0",
  "userAgent": "aws-sdk-rust/0.55.2 os/linux lang/rust/1.69.0",
  "requestParameters": {
    "maxResults": 10
  },
  "responseElements": null,
  "requestID": "5ef238db-9f87-4f37-ab7b-6cf0ba5df891",
  "eventID": "b0430fb0-12c3-4cca-8d05-84c37f99c51f",
  "readOnly": true,
  "eventType": "AwsApiCall",
  "managementEvent": true,
  "recipientAccountId": "123456789012",
  "eventCategory": "Management"
}
```

DeletePolicyStore

The following example shows a CloudTrail log entry for a `DeletePolicyStore` API call:


```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "EXAMPLE_PRINCIPAL_ID",
    "arn": "arn:aws:iam::123456789012:role/ExampleRole",
    "accountId": "123456789012",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE"
  },
  "eventTime": "2023-05-22T07:43:32Z",
  "eventSource": "verifiedpermissions.amazonaws.com",
  "eventName": "DeletePolicyStore",
  "awsRegion": "us-west-2",
  "sourceIPAddress": "203.0.113.0",
  "userAgent": "aws-sdk-rust/0.55.2 os/linux lang/rust/1.69.0",
  "requestParameters": {
    "policyStoreId": "PSEXAMPLEabcdefg111111"
  },
  "responseElements": null,
  "requestID": "1368e8f9-130d-45a5-b96d-99097ca3077f",
  "eventID": "ac482022-b2f6-4069-879a-dd509123d8d7",
  "readOnly": false,
  "resources": [
    {
      "accountId": "123456789012",
      "type": "AWS::VerifiedPermissions::PolicyStore",
      "arn": "arn:aws:verifiedpermissions::123456789012:policy-store/
PSEXAMPLEabcdefg111111"
    }
  ],
  "eventType": "AwsApiCall",
  "managementEvent": true,
  "recipientAccountId": "123456789012",
  "eventCategory": "Management"
}
```


PutSchema

The following example shows a CloudTrail log entry for a `PutSchema` API call:

```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "EXAMPLE_PRINCIPAL_ID",
    "arn": "arn:aws:iam::123456789012:role/ExampleRole",
    "accountId": "123456789012",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE"
  },
  "eventTime": "2023-05-16T12:58:57Z",
  "eventSource": "verifiedpermissions.amazonaws.com",
  "eventName": "PutSchema",
  "awsRegion": "us-west-2",
  "sourceIPAddress": "203.0.113.0",
  "userAgent": "aws-sdk-rust/0.55.2 os/linux lang/rust/1.69.0",
  "requestParameters": {
    "policyStoreId": "PSEXAMPLEabcdefg111111"
  },
  "responseElements": {
    "lastUpdatedDate": "2023-05-16T12:58:57.513442Z",
    "namespaces": "[some_namespace]",
    "createdDate": "2023-05-16T12:58:57.513442Z",
    "policyStoreId": "PSEXAMPLEabcdefg111111",
  },
  "requestID": "631fbfa1-a959-4988-b9f8-f1a43ff5df0d",
  "eventID": "7cd0c677-733f-4602-bc03-248bae581fe5",
  "readOnly": false,
  "resources": [
    {
      "accountId": "123456789012",
      "type": "AWS::VerifiedPermissions::PolicyStore",
      "ARN": "arn:aws:verifiedpermissions::123456789012:policy-store/
PSEXAMPLEabcdefg111111"
    }
  ],
  "eventType": "AwsApiCall",
  "managementEvent": true,
  "recipientAccountId": "123456789012",
  "eventCategory": "Management"
}
```

GetSchema

The following example shows a CloudTrail log entry for a `GetSchema` API call:


```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "EXAMPLE_PRINCIPAL_ID",
    "arn": "arn:aws:iam::222222222222:role/ExampleRole",
    "accountId": "222222222222",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE"
  },
  "eventTime": "2023-05-25T01:12:07Z",
  "eventSource": "verifiedpermissions.amazonaws.com",
  "eventName": "GetSchema",
  "awsRegion": "us-west-2",
  "sourceIPAddress": "203.0.113.0",
  "userAgent": "aws-sdk-rust/0.55.2 os/linux lang/rust/1.69.0",
  "requestParameters": {
    "policyStoreId": "PSEXAMPLEabcdefg111111"
  },
  "responseElements": null,
  "requestID": "a1f4d4cd-6156-480a-a9b8-e85a71dcc7c2",
  "eventID": "0b3b8e3d-155c-46f3-a303-7e9e8b5f606b",
  "readOnly": true,
  "resources": [
    {
      "accountId": "222222222222",
      "type": "AWS::VerifiedPermissions::PolicyStore",
      "ARN": "arn:aws:verifiedpermissions::222222222222:policy-store/
PSEXAMPLEabcdefg111111"
    }
  ],
  "eventType": "AwsApiCall",
  "managementEvent": true,
  "recipientAccountId": "222222222222",
  "eventCategory": "Management"
}
```


CreatePolicyTemplate

The following example shows a CloudTrail log entry for a `CreatePolicyTemplate` API call:

```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "EXAMPLE_PRINCIPAL_ID",
    "arn": "arn:aws:iam::123456789012:role/ExampleRole",
    "accountId": "123456789012",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE"
  },
  "eventTime": "2023-05-16T13:00:24Z",
  "eventSource": "verifiedpermissions.amazonaws.com",
  "eventName": "CreatePolicyTemplate",
  "awsRegion": "us-west-2",
  "sourceIPAddress": "203.0.113.0",
  "userAgent": "aws-sdk-rust/0.55.2 os/linux lang/rust/1.69.0",
  "requestParameters": {
    "policyStoreId": "PSEXAMPLEabcdefg111111"
  },
  "responseElements": {
    "lastUpdatedDate": "2023-05-16T13:00:23.444404Z",
    "createdDate": "2023-05-16T13:00:23.444404Z",
    "policyTemplateId": "PTEXAMPLEabcdefg111111",
    "policyStoreId": "PSEXAMPLEabcdefg111111",
  },
  "requestID": "73953bda-af5e-4854-afe2-7660b492a6d0",
  "eventID": "7425de77-ed84-4f91-a4b9-b669181cc57b",
  "readOnly": false,
  "resources": [
    {
      "accountId": "123456789012",
      "type": "AWS::VerifiedPermissions::PolicyStore",
      "arn": "arn:aws:verifiedpermissions::123456789012:policy-store/
PSEXAMPLEabcdefg111111"
    }
  ],
  "eventType": "AwsApiCall",
  "managementEvent": true,
  "recipientAccountId": "123456789012",
  "eventCategory": "Management"
}
```

DeletePolicyTemplate

The following example shows a CloudTrail log entry for a `DeletePolicyTemplate` API call:


```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "EXAMPLE_PRINCIPAL_ID",
    "arn": "arn:aws:iam::222222222222:role/ExampleRole",
    "accountId": "222222222222",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE"
  },
  "eventTime": "2023-05-25T01:11:48Z",
  "eventSource": "verifiedpermissions.amazonaws.com",
  "eventName": "DeletePolicyTemplate",
  "awsRegion": "us-west-2",
  "sourceIPAddress": "203.0.113.0",
  "userAgent": "aws-sdk-rust/0.55.2 os/linux lang/rust/1.69.0",
  "requestParameters": {
    "policyStoreId": "PSEXAMPLEabcdefg111111",
    "policyTemplateId": "PTEXAMPLEabcdefg111111"
  },
  "responseElements": null,
  "requestID": "5ff0f22e-6bbd-4b85-a400-4fb74aa05dc6",
  "eventID": "c0e0c689-369e-4e95-a9cd-8de113d47ffa",
  "readOnly": false,
  "resources": [
    {
      "accountId": "222222222222",
      "type": "AWS::VerifiedPermissions::PolicyStore",
      "ARN": "arn:aws:verifiedpermissions::222222222222:policy-store/
PSEXAMPLEabcdefg111111"
    }
  ],
  "eventType": "AwsApiCall",
  "managementEvent": true,
  "recipientAccountId": "222222222222",
  "eventCategory": "Management"
}
```

CreatePolicy

The following example shows a CloudTrail log entry for a `CreatePolicy` API call:


```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "EXAMPLE_PRINCIPAL_ID",
    "arn": "arn:aws:iam::123456789012:role/ExampleRole",
    "accountId": "123456789012",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE"
  },
  "eventTime": "2023-05-22T07:42:30Z",
  "eventSource": "verifiedpermissions.amazonaws.com",
  "eventName": "CreatePolicy",
  "awsRegion": "us-west-2",
  "sourceIPAddress": "203.0.113.0",
  "userAgent": "aws-sdk-rust/0.55.2 os/linux lang/rust/1.69.0",
  "requestParameters": {
    "clientToken": "a1b2c3d4-e5f6-a1b2-c3d4-TOKEN1111111",
    "policyStoreId": "PSEXAMPLEabcdefg111111"
  },
  "responseElements": {
    "policyStoreId": "PSEXAMPLEabcdefg111111",
    "policyId": "SPEXAMPLEabcdefg111111",
    "policyType": "STATIC",
    "principal": {
        "entityType": "PhotoApp::Role",
        "entityId": "PhotoJudge"
    },
    "resource": {
        "entityType": "PhotoApp::Application",
        "entityId": "PhotoApp"
    },
    "lastUpdatedDate": "2023-05-22T07:42:30.70852Z",
    "createdDate": "2023-05-22T07:42:30.70852Z"
  },
  "requestID": "93ffa151-3841-4960-9af6-30a7f817ef93",
  "eventID": "30ab405f-3dff-43ff-8af9-f513829e8bde",
  "readOnly": false,
  "resources": [
    {
      "accountId": "123456789012",
```

```
      "type": "AWS::VerifiedPermissions::PolicyStore",
      "arn": "arn:aws:verifiedpermissions::123456789012:policy-store/
PSEXAMPLEabcdefg111111"
    }
  ],
  "eventType": "AwsApiCall",
  "managementEvent": true,
  "recipientAccountId": "123456789012",
  "eventCategory": "Management"
}
```

GetPolicy

The following example shows a CloudTrail log entry for a `GetPolicy` API call:

```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "EXAMPLE_PRINCIPAL_ID",
    "arn": "arn:aws:iam::123456789012:role/ExampleRole",
    "accountId": "123456789012",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE"
  },
  "eventTime": "2023-05-22T07:43:29Z",
  "eventSource": "verifiedpermissions.amazonaws.com",
  "eventName": "GetPolicy",
  "awsRegion": "us-west-2",
  "sourceIPAddress": "203.0.113.0",
  "userAgent": "aws-sdk-rust/0.55.2 os/linux lang/rust/1.69.0",
  "requestParameters": {
    "policyStoreId": "PSEXAMPLEabcdefg111111",
    "policyId": "SPEXAMPLEabcdefg111111"
  },
  "responseElements": null,
  "requestID": "23022a9e-2f5c-4dac-b653-59e6987f2fac",
  "eventID": "9b4d5037-bafa-4d57-b197-f46af83fc684",
  "readOnly": true,
  "resources": [
    {
      "accountId": "123456789012",
      "type": "AWS::VerifiedPermissions::PolicyStore",
```

```
      "arn": "arn:aws:verifiedpermissions::123456789012:policy-store/
PSEXAMPLEabcdefg111111"
    }
  ],
  "eventType": "AwsApiCall",
  "managementEvent": true,
  "recipientAccountId": "123456789012",
  "eventCategory": "Management"
}
```

CreateIdentitySource

The following example shows a CloudTrail log entry for a `CreateIdentitySource` API call:

```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "EXAMPLE_PRINCIPAL_ID",
    "arn": "arn:aws:iam::333333333333:role/ExampleRole",
    "accountId": "333333333333",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE"
  },
  "eventTime": "2023-05-19T01:27:44Z",
  "eventSource": "verifiedpermissions.amazonaws.com",
  "eventName": "CreateIdentitySource",
  "awsRegion": "us-west-2",
  "sourceIPAddress": "203.0.113.0",
  "userAgent": "aws-sdk-rust/0.55.2 os/linux lang/rust/1.69.0",
  "requestParameters": {
    "clientToken": "a1b2c3d4-e5f6-a1b2-c3d4-TOKEN1111111",
    "configuration": {
      "cognitoUserPoolConfiguration": {
        "userPoolArn": "arn:aws:cognito-idp:000011112222:us-east-1:userpool/us-
east-1_aaaaaaaaaa"
      }
    },
    "policyStoreId": "PSEXAMPLEabcdefg111111",
    "principalEntityType": "User"
  },
  "responseElements": {
    "createdDate": "2023-07-14T15:05:01.599534Z",
    "identitySourceId": "ISEXAMPLEabcdefg111111",
```

```
    "lastUpdatedDate": "2023-07-14T15:05:01.599534Z",
    "policyStoreId": "PSEXAMPLEabcdefg111111"
    },
  "requestID": "afcc1e67-d5a4-4a9b-a74c-cdc2f719391c",
  "eventID": "f13a41dc-4496-4517-aeb8-a389eb379860",
  "readOnly": false,
  "resources": [
    {
      "accountId": "333333333333",
      "type": "AWS::VerifiedPermissions::PolicyStore",
      "arn": "arn:aws:verifiedpermissions::333333333333:policy-store/
PSEXAMPLEabcdefg111111"
    }
  ],
  "eventType": "AwsApiCall",
  "managementEvent": true,
  "recipientAccountId": "333333333333",
  "eventCategory": "Management"
}
```

GetIdentitySource

The following example shows a CloudTrail log entry for a `GetIdentitySource` API call:

```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "EXAMPLE_PRINCIPAL_ID",
    "arn": "arn:aws:iam::333333333333:role/ExampleRole",
    "accountId": "333333333333",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE"
  },
  "eventTime": "2023-05-24T19:55:31Z",
  "eventSource": "verifiedpermissions.amazonaws.com",
  "eventName": "GetIdentitySource",
  "awsRegion": "us-west-2",
  "sourceIPAddress": "203.0.113.0",
  "userAgent": "aws-sdk-rust/0.55.2 os/linux lang/rust/1.69.0",
  "requestParameters": {
    "identitySourceId": "ISEXAMPLEabcdefg111111",
    "policyStoreId": "PSEXAMPLEabcdefg111111"
  },
```

```
  "responseElements": null,
  "requestID": "7a6ecf79-c489-4516-bb57-9ded970279c9",
  "eventID": "fa158e6c-f705-4a15-a731-2cdb4bd9a427",
  "readOnly": true,
  "resources": [
    {
      "accountId": "333333333333",
      "type": "AWS::VerifiedPermissions::PolicyStore",
      "arn": "arn:aws:verifiedpermissions::333333333333:policy-store/
PSEXAMPLEabcdefg111111"
    }
  ],
  "eventType": "AwsApiCall",
  "managementEvent": true,
  "recipientAccountId": "333333333333",
  "eventCategory": "Management"
}
```

ListIdentitySources

The following example shows a CloudTrail log entry for a `ListIdentitySources` API call:

```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "EXAMPLE_PRINCIPAL_ID",
    "arn": "arn:aws:iam::333333333333:role/ExampleRole",
    "accountId": "333333333333",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE"
  },
  "eventTime": "2023-05-24T20:05:32Z",
  "eventSource": "verifiedpermissions.amazonaws.com",
  "eventName": "ListIdentitySources",
  "awsRegion": "us-west-2",
  "sourceIPAddress": "203.0.113.0",
  "userAgent": "aws-sdk-rust/0.55.2 os/linux lang/rust/1.69.0",
  "requestParameters": {
    "policyStoreId": "PSEXAMPLEabcdefg111111"
  },
  "responseElements": null,
  "requestID": "95d2a7bc-7e9a-4efe-918e-97e558aacaf7",
  "eventID": "d3dc53f6-1432-40c8-9d1d-b9eeb75c6193",
```

```
  "readOnly": true,
  "resources": [
    {
      "accountId": "333333333333",
      "type": "AWS::VerifiedPermissions::PolicyStore",
      "arn": "arn:aws:verifiedpermissions::333333333333:policy-store/
PSEXAMPLEabcdefg111111"
    }
  ],
  "eventType": "AwsApiCall",
  "managementEvent": true,
  "recipientAccountId": "333333333333",
  "eventCategory": "Management"
}
```

DeleteIdentitySource

The following example shows a CloudTrail log entry for a `DeleteIdentitySource` API call:

```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "EXAMPLE_PRINCIPAL_ID",
    "arn": "arn:aws:iam::333333333333:role/ExampleRole",
    "accountId": "333333333333",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE"
  },
  "eventTime": "2023-05-24T19:55:32Z",
  "eventSource": "verifiedpermissions.amazonaws.com",
  "eventName": "DeleteIdentitySource",
  "awsRegion": "us-west-2",
  "sourceIPAddress": "203.0.113.0",
  "userAgent": "aws-sdk-rust/0.55.2 os/linux lang/rust/1.69.0",
  "requestParameters": {
    "identitySourceId": "ISEXAMPLEabcdefg111111",
    "policyStoreId": "PSEXAMPLEabcdefg111111"
  },
  "responseElements": null,
  "requestID": "d554d964-0957-4834-a421-c417bd293086",
  "eventID": "fe4d867c-88ee-4e5d-8d30-2fbc208c9260",
  "readOnly": false,
  "resources": [
```

```
    {
      "accountId": "333333333333",
      "type": "AWS::VerifiedPermissions::PolicyStore",
      "arn": "arn:aws:verifiedpermissions::333333333333:policy-store/
PSEXAMPLEabcdefg111111"
    }
  ],
  "eventType": "AwsApiCall",
  "managementEvent": true,
  "recipientAccountId": "333333333333",
  "eventCategory": "Management"
}
```

Creating Amazon Verified Permissions resources with AWS CloudFormation

Amazon Verified Permissions is integrated with AWS CloudFormation, a service that helps you to model and set up your AWS resources so that you can spend less time creating and managing your resources and infrastructure. You create a template that describes all the AWS resources that you want (such as policy stores), and CloudFormation provisions and configures those resources for you.

When you use CloudFormation, you can reuse your template to set up your Verified Permissions resources consistently and repeatedly. Describe your resources once, and then provision the same resources over and over in multiple AWS accounts and Regions.

> **Important**
>
> Amazon Cognito Identity is not available in all of the same AWS Regions as Amazon Verified Permissions. If you receive an error from CloudFormation regarding Amazon Cognito Identity, such as `Unrecognized resource types: AWS::Cognito::UserPool, AWS::Cognito::UserPoolClient`, we recommend that you create the Amazon Cognito user pool and client in the geographically closest AWS Region where Amazon Cognito Identity is available. Use this newly created user pool when creating the Verified Permissions identity source.

Verified Permissions and CloudFormation templates

To provision and configure resources for Verified Permissions and related services, you must understand CloudFormation templates. Templates are formatted text files in JSON or YAML. These templates describe the resources that you want to provision in your CloudFormation stacks. If you're unfamiliar with JSON or YAML, you can use CloudFormation Designer to help you get started with CloudFormation templates. For more information, see What is CloudFormation Designer? in the *AWS CloudFormation User Guide*.

Verified Permissions supports creating identity sources, policies, policy stores, policy templates, and policy store aliases in CloudFormation. For more information, including examples of JSON and YAML templates for Verified Permissions resources, see the Amazon Verified Permissions resource type reference in the *AWS CloudFormation User Guide*.

AWS CDK constructs

The AWS Cloud Development Kit (AWS CDK) is an open-source software development framework for defining cloud infrastructure in code and provisioning it through CloudFormation. Constructs, or reusable cloud components, can be used to create CloudFormation templates. These templates can then be used to deploy your cloud infrastructure.

To learn more and download AWS CDK, see AWS Cloud Development Kit.

The following are links to documentation for Verified Permissions AWS CDK resources, such as constructs.

- Amazon Verified Permissions L2 CDK Construct

Learn more about CloudFormation

To learn more about CloudFormation, see the following resources:

- AWS CloudFormation
- AWS CloudFormation User Guide
- CloudFormation API Reference
- AWS CloudFormation Command Line Interface User Guide

Access Amazon Verified Permissions using AWS PrivateLink

You can use AWS PrivateLink to create a private connection between your VPC and Amazon Verified Permissions. You can access Verified Permissions as if it were in your VPC, without the use of an internet gateway, NAT device, VPN connection, or Direct Connect connection. Instances in your VPC don't need public IP addresses to access Verified Permissions.

You establish this private connection by creating an *interface endpoint*, powered by AWS PrivateLink. We create an endpoint network interface in each subnet that you enable for the interface endpoint. These are requester-managed network interfaces that serve as the entry point for traffic destined for Verified Permissions.

For more information, see Access AWS services through AWS PrivateLink in the *AWS PrivateLink Guide*.

Considerations for Verified Permissions

Before you set up an interface endpoint for Verified Permissions, review Considerations in the *AWS PrivateLink Guide*.

Verified Permissions supports making calls to all of its API actions through the interface endpoint.

VPC endpoint policies are not supported for Verified Permissions. By default, full access to Verified Permissions is allowed through the interface endpoint. Alternatively, you can associate a security group with the endpoint network interfaces to control traffic to Verified Permissions through the interface endpoint.

Create an interface endpoint for Verified Permissions

You can create an interface endpoint for Verified Permissions using either the Amazon VPC console or the AWS Command Line Interface (AWS CLI). For more information, see Create an interface endpoint in the *AWS PrivateLink Guide*.

Create an interface endpoint for Verified Permissions using the following service name:

```
com.amazonaws.region.verifiedpermissions
```

If you enable private DNS for the interface endpoint, you can make API requests to Verified Permissions using its default Regional DNS name. For example, `verifiedpermissions.us-east-1.amazonaws.com`.

Create an endpoint policy for your interface endpoint

An endpoint policy is an IAM resource that you can attach to an interface endpoint. The default endpoint policy allows full access to Verified Permissions through the interface endpoint. To control the access allowed to Verified Permissions from your VPC, attach a custom endpoint policy to the interface endpoint.

An endpoint policy specifies the following information:

- The principals that can perform actions (AWS accounts, IAM users, and IAM roles).
- The actions that can be performed.
- The resources on which the actions can be performed.

For more information, see Control access to services using endpoint policies in the *AWS PrivateLink Guide*.

Example: VPC endpoint policy for Verified Permissions actions

The following is an example of a custom endpoint policy. When you attach this policy to your interface endpoint, it grants access to the listed Verified Permissions actions for all principals on all resources.

```
{
   "Statement": [
      {
         "Principal": "*",
         "Effect": "Allow",
         "Action": [
            "verifiedpermissions:IsAuthorized",
            "verifiedpermissions:IsAuthorizedWithToken",
            "verifiedpermissions:GetPolicy"
         ],
         "Resource":"*"
      }
   ]
```

```
}
```

Quotas for Amazon Verified Permissions

Your AWS account has default quotas, formerly referred to as limits, for each AWS service. Unless otherwise noted, each quota is Region-specific. You can request increases for some quotas, and other quotas cannot be increased.

To view the quotas for Verified Permissions, open the Service Quotas console. In the navigation pane, choose **AWS services** and select **Verified Permissions**.

To request a quota increase, see Requesting a Quota Increase in the *Service Quotas User Guide*. If the quota is not yet available in Service Quotas, use the limit increase form.

Your AWS account has the following quotas related to Verified Permissions.

Topics

Quotas for resources

Name	Default	Adjuse	Description
Policy stores per Region per account	Each supported Region: 30,000	Yes	The maximum number of policy stores.
Policy templates per policy store	Each supported Region: 40	Yes	The maximum number of policy templates in a policy store.
Identity sources per policy store	1	No	The maximum number of identity sources that you can define for a policy store.

Name	Default	Adjustable	Description
Policy store aliases per policy store	10	Yes	The maximum number of policy store aliases that you can associate with a single policy store.
Authorization request size[1]	1 MB	No	The maximum size of an authorization request.
Policy size	10,000 bytes	Yes	The maximum size of an individual policy.
Schema size	100,000 bytes	Yes	The maximum size of the schema of a policy store.
Namespaces per policy store schema	100	Yes	The maximum number of namespaces that you can define in a policy store schema.
Policy size per resource	200,000 bytes[2]	Yes	The maximum size of all policies that reference a specific resource.

[1] The quota for an authorization request is the same for both IsAuthorized and IsAuthorizedWithToken.

[2] The default limit for the total size of all the policies that reference a single resource is 200,000 bytes. Policies that don't specify a resource share a separate total of 200,000 bytes for the `"unspecified"` resource. For details about how Verified Permissions calculates this quota, examples, and strategies for staying within it, see Policy size per resource.

Quotas for hierarchies

> **Note**
>
> Each transitive parent quota applies to each entity individually and counts both direct and indirect (transitive) parents toward the total. For example, if the limit of *Transitive parents per principal* is 100, a principal can have 100 direct parent groups with no nested groups. Alternatively, a principal can have 10 direct parent groups that each belong to 9 additional parent groups. Any combination that totals 100 parents for that principal is valid.

Name	Default	Adjustable	Description
Transitive parents per principal	100	No	The maximum number of transitive parents for each principal.
Transitive parents per action	100	No	The maximum number of transitive parents for each action.
Transitive parents per resource	100	No	The maximum number of transitive parents for each resource.

The diagram below illustrates how transitive parents can be defined for an entity (principal, action, or resource).

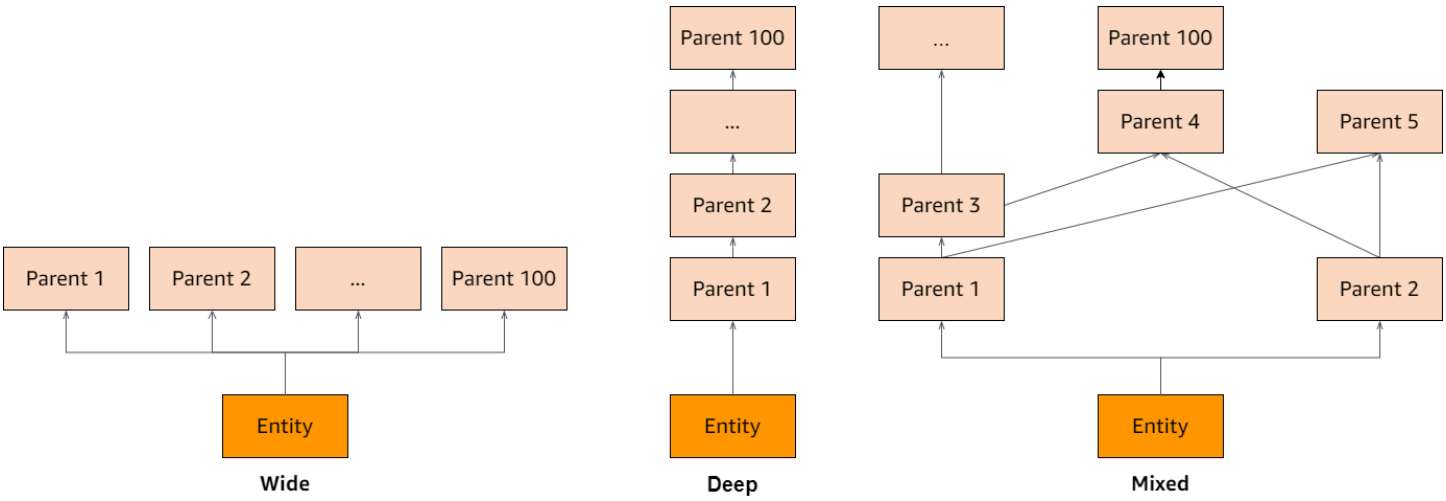

Quotas for operations per second

Verified Permissions throttles requests to service endpoints in an AWS Region when application requests exceed the quota for an API operation. Verified Permissions might return an exception when you exceed the quota in requests per second, or you attempt simultaneous write operations. You can view your current RPS quotas in Service Quotas. To prevent applications from exceeding the quota for an operation, you must optimize them for retries and exponential backoff. For more information, see Retry with backoff pattern and Managing and monitoring API throttling in your workloads.

Name	Default	Adjustable	Description
BatchGetPolicy requests per second per Region per policy store	Each supported Region: 10	Yes	The maximum number of BatchGetPolicy requests per second per policy store.
BatchIsAuthorized requests per second per Region per policy store	Each supported Region: 30	Yes	The maximum number of BatchIsAuthorized requests per second per policy store.
BatchIsAuthorizedWithToken requests per second per Region per policy store	Each supported Region: 30	Yes	The maximum number of BatchIsAuthorizedW

Name	Default	Adjustable	Description
			ithToken requests per second per policy store.
CreateIdentitySource requests per second per Region per policy store	Each supported Region: 1	Yes	The maximum number of CreateIdentitySource requests per second per policy store.
CreatePolicy requests per second per Region per policy store	Each supported Region: 10	Yes	The maximum number of CreatePolicy requests per second per policy store.
CreatePolicyStore requests per second per Region per account	Each supported Region: 1	No	The maximum number of CreatePolicyStore requests per second.
CreatePolicyTemplate requests per second per Region per policy store	Each supported Region: 10	Yes	The maximum number of CreatePolicyTemplate requests per second per policy store.
DeleteIdentitySource requests per second per Region per policy store	Each supported Region: 1	Yes	The maximum number of DeleteIdentitySource requests per second per policy store.
DeletePolicy requests per second per Region per policy store	Each supported Region: 10	Yes	The maximum number of DeletePolicy requests per second per policy store.
DeletePolicyStore requests per second per Region per account	Each supported Region: 1	No	The maximum number of DeletePolicyStore requests per second.

Name	Default	Adjustable	Description
DeletePolicyTemplate requests per second per Region per policy store	Each supported Region: 10	Yes	The maximum number of DeletePolicyTemplate requests per second per policy store.
GetIdentitySource requests per second per Region per policy store	Each supported Region: 10	Yes	The maximum number of GetIdentitySource requests per second per policy store.
GetPolicy requests per second per Region per policy store	Each supported Region: 10	Yes	The maximum number of GetPolicy requests per second per policy store.
GetPolicyStore requests per second per Region per account	Each supported Region: 10	Yes	The maximum number of GetPolicyStore requests per second.
GetPolicyTemplate requests per second per Region per policy store	Each supported Region: 10	Yes	The maximum number of GetPolicyTemplate requests per second per policy store.
GetSchema requests per second per Region per policy store	Each supported Region: 10	Yes	The maximum number of GetSchema requests per second per policy store.
IsAuthorized requests per second per Region per policy store	Each supported Region: 200	Yes	The maximum number of IsAuthorized requests per second per policy store.
IsAuthorizedWithToken requests per second per Region per policy store	Each supported Region: 200	Yes	The maximum number of IsAuthorizedWithToken requests per second per policy store.

Name	Default	Adjustable	Description
ListIdentitySources requests per second per Region per policy store	Each supported Region: 10	Yes	The maximum number of ListIdentitySources requests per second per policy store.
ListPolicies requests per second per Region per policy store	Each supported Region: 10	Yes	The maximum number of ListPolicies requests per second per policy store.
ListPolicyStores requests per second per Region per account	Each supported Region: 10	Yes	The maximum number of ListPolicyStores requests per second.
ListPolicyTemplates requests per second per Region per policy store	Each supported Region: 10	Yes	The maximum number of ListPolicyTemplates requests per second per policy store.
PutSchema requests per second per Region per policy store	Each supported Region: 10	Yes	The maximum number of PutSchema requests per second per policy store.
UpdateIdentitySource requests per second per Region per policy store	Each supported Region: 1	Yes	The maximum number of UpdateIdentitySource requests per second per policy store.
UpdatePolicy requests per second per Region per policy store	Each supported Region: 10	Yes	The maximum number of UpdatePolicy requests per second per policy store.
UpdatePolicyStore requests per second per Region per account	Each supported Region: 10	No	The maximum number of UpdatePolicyStore requests per second.

Name	Default	Adjustable	Description
UpdatePolicyTemplate requests per second per Region per policy store	Each supported Region: 10	Yes	The maximum number of UpdatePolicyTemplate requests per second per policy store.

Amazon Verified Permissions and Cedar policy language terms and concepts

You should understand the following concepts to use Amazon Verified Permissions.

Verified Permissions concepts

- Authorization model
- Authorization request
- Authorization response
- Considered policies
- Context data
- Determining policies
- Entity data
- Permissions, authorization, and principals
- Policy enforcement
- Policy store
- Policy store alias
- Policy name
- Policy template name
- Satisfied policies
- Differences between Amazon Verified Permissions and the Cedar policy language

Cedar policy language concepts

- Authorization
- Entity
- Groups and hierarchies
- Namespaces
- Policy
- Policy template

- Schema

Authorization model

The *authorization model* describes the scope of the authorization requests made by the application and the basis for evaluating those requests. It is defined in terms of the different types of resources, the actions taken on those resources, and the types principals that take those actions. It also considers the context in which those actions are being taken.

Role-based Access Control (RBAC) is an evaluation basis in which roles are defined and associated with a set of permissions. These roles can then be assigned to one or more identities. The assigned identity acquires the permissions associated with the role. If the permissions associated with the role are modified, then the modification automatically impacts any identity to which the role has been assigned. Cedar can support RBAC decisions through the use of principal groups.

Attribute-based Access Control (ABAC) is an evaluation basis in which the permissions associated with an identity are determined by attributes of that identity. Cedar can support ABAC decisions through the use of policy conditions that reference attributes of the principal.

The Cedar policy language enables the combination of RBAC and ABAC in a single policy by allowing permissions to be defined for a group of users, which have attribute-based conditions.

Authorization request

An *authorization request* is a request made of Verified Permissions by an application to evaluate a set of policies in order to determine whether a principal may perform an action on a resource for a given context.

Authorization response

The *authorization response* is the response to the authorization request. It includes an allow or deny decision, plus additional information, such as the IDs of the determining policies.

Considered policies

Considered policies are the full set of policies that are selected by Verified Permissions for inclusion when evaluating an authorization request.

Context data

Context data are attribute values that provide additional information to be evaluated.

Determining policies

Determining policies are the policies that determine the authorization response. For example, if there are two satisfied policies, where one is a deny and the other is an allow, then the deny policy will be the determining policy. If there are multiple satisfied permit policies and no satisfied forbid policies, then there are multiple determining policies. In the case that no policies match and the response is deny, there are no determining policies.

Entity data

Entity data are data about the principal, action, and resource. Entity data relevant for policy evaluation are group membership all the way up the entity hierarchy and attribute values of the principal and resource.

Permissions, authorization, and principals

Verified Permissions manages fine-grained *permissions* and *authorization* within custom applications that you build.

A *principal* is user of an application, either human or machine, that has an identity bound to an identifier such as a username or machine ID. The process of authentication determines whether the principal is truly the identity they claim to be.

Associated with that identity are a set of application *permissions* that determine what that principal is permitted to do within that application. *Authorization* is the process of assessing those permissions to determine whether a principal is permitted to perform a particular action in the application. These permissions can be expressed as policies.

Policy enforcement

Policy enforcement is the process of enforcing the evaluation decision within the application outside of Verified Permissions. If Verified Permissions evaluation returns a deny, then enforcement would ensure that the principal was prevented from accessing the resource.

Policy store

A *policy store* is a container for policies and templates. Each store contains a schema that is used to validate policies added to the store. By default, each application has its own policy store, but multiple applications can share a single policy store. When an application makes an authorization request, it identifies the policy store used to evaluate that request. Policy stores provide a way to isolate a set of policies, and can therefore be used in a multi-tenant application to contain the schemas and policies for each tenant. A single application can have separate policy stores for each tenant.

When evaluating an authorization request, Verified Permissions only considers the subset of the policies in the policy store that are relevant to the request. Relevance is determined based on the *scope* of the policy. The scope identifies the specific principal and resource to which the policy applies, and the actions that the principal can perform on the resource. Defining the scope helps improve performance by narrowing the set of considered policies.

Policy store alias

A *policy store alias* is a friendly name for a policy store. You can use a policy store alias to identify a policy store in any Verified Permissions operation that accepts a `policyStoreId` parameter. Policy store aliases are independent AWS resources with their own ARNs. Each alias is associated with one policy store at a time, and multiple aliases can be associated with the same policy store. For more information, see Amazon Verified Permissions policy store aliases.

Policy name

A *policy name* is an optional friendly name for a policy. Policy names must be unique for all policies within the policy store and prefixed with `name/`. You can use a policy name in place of the policy ID in control plane operations that accept a `policyId` parameter. Names can be set when creating or updating a policy. Only `GetPolicy` and `ListPolicies` return the name in the output.

Policy template name

A *policy template name* is an optional friendly name for a policy template. Policy template names must be unique for all policy templates within the policy store and prefixed with `name/`. You can use a policy template name in place of the policy template ID in control plane operations that

accept a `policyTemplateId` parameter. Names can be set when creating or updating a policy template. Only `GetPolicyTemplate` and `ListPolicyTemplates` return the name in the output.

Satisfied policies

Satisfied policies are the policies that match the parameters of the authorization request.

Differences between Amazon Verified Permissions and the Cedar policy language

Amazon Verified Permissions uses the Cedar policy language engine to perform its authorization tasks. However, there are some differences between the native Cedar implementation and the implementation of Cedar found in Verified Permissions. This topic identifies those differences.

> **Note**
>
> As of May 2026, Verified Permissions has aligned with Cedar and now supports multiple namespaces.

Policy template support

Both Verified Permissions and Cedar allow placeholders in the scope for only the `principal` and `resource`. However, Verified Permissions also requires that neither the `principal` and `resource` are unconstrained.

The following policy is valid in Cedar but is rejected by Verified Permissions because the `principal` is unconstrained.

```
permit(principal, action == Action::"view", resource == ?resource);
```

Both of the following examples are valid in both Cedar and Verified Permissions because both the `principal` and `resource` have constraints.

```
permit(principal == User::"alice", action == Action::"view", resource == ?resource);
```

```
permit(principal == ?principal, action == Action::"a", resource in ?resource);
```

Schema support

Verified Permissions requires all schema JSON key names to be non-empty strings. Cedar allows empty strings in a few cases, such as for properties or namespaces.

Action groups definition

The Cedar authorization methods require a list of the entities to be considered when evaluating an authorization request against the policies.

You can define the actions and action groups used by your application in the schema. However, Cedar doesn't include the schema as part of an evaluation request. Instead, Cedar uses the schema only to validate the policies and policy templates that you submit. Because Cedar doesn't reference the schema during evaluation requests, even if you defined action groups in the schema, you must also include the list of any action groups as part of the entities list you must pass to the authorization API operations.

Verified Permissions does this for you. Any action groups that you define in your schema are automatically appended to the entities list that you pass to as a parameter to the `IsAuthorized` or `IsAuthorizedWithToken` operations.

Entity formatting

The JSON formatting of entities in Verified Permissions using the `entityList` parameter differs from Cedar in the following ways:

- In Verified Permissions, a JSON object must have all of its key-value pairs wrapped in a JSON object with the name of `Record`.
- A JSON list in Verified Permissions must be wrapped in a JSON key-value pair where the key name is `Set` and the value is the original JSON list from Cedar.
- For `String`, `Long`, and `Boolean` type names, each key-value pair from Cedar is replaced by a JSON object in Verified Permissions. The name of the object is the original key name. Inside the JSON object, there is one key-value pair where the key name is the type name of the scalar value (`String`, `Long`, or `Boolean`) and the value is the value from the Cedar entity.
- The syntax formatting of Cedar entities and Verified Permissions entities differs in the following ways:

Cedar format	Verified Permissions format
uid	Identifier
type	EntityType
id	EntityId
attrs	Attributes
parents	Parents

Example- Lists

The following examples show how a list of entities is expressed in Cedar and Verified Permissions, respectively.

Cedar

```
[
  {
    "number": 1
  },
  {
    "sentence": "Here is an example sentence"
  },
  {
    "Question": false
  }
]
```

Verified Permissions

```
{
  "Set": [
    {
      "Record": {
        "number": {
          "Long": 1
        }
```

```
      }
    },
    {
      "Record": {
        "sentence": {
          "String": "Here is an example sentence"
        }
      }
    },
    {
      "Record": {
        "question": {
          "Boolean": false
        }
      }
    }
  ]
}
```

Example- Policy evaluation

The following examples shows how entities are formatted for evaluating a policy in an authorization request in Cedar and Verified Permissions, respectively.

Cedar

```
[
    {
        "uid": {
            "type": "PhotoApp::User",
            "id": "alice"
        },
        "attrs": {
            "age": 25,
            "name": "alice",
            "userId": "123456789012"
        },
        "parents": [
            {
                "type": "PhotoApp::UserGroup",
                "id": "alice_friends"
            },
```

```
            {
                "type": "PhotoApp::UserGroup",
                "id": "AVTeam"
            }
        ]
    },
    {
        "uid": {
            "type": "PhotoApp::Photo",
            "id": "vacationPhoto.jpg"
        },
        "attrs": {
            "private": false,
            "account": {
                "__entity": {
                    "type": "PhotoApp::Account",
                    "id": "ahmad"
                }
            }
        },
        "parents": []
    },
    {
        "uid": {
            "type": "PhotoApp::UserGroup",
            "id": "alice_friends"
        },
        "attrs": {},
        "parents": []
    },
    {
        "uid": {
            "type": "PhotoApp::UserGroup",
            "id": "AVTeam"
        },
        "attrs": {},
        "parents": []
    }
]
```

Verified Permissions

```
[
```

```
    {
        "Identifier": {
            "EntityType": "PhotoApp::User",
            "EntityId": "alice"
        },
        "Attributes": {
            "age": {
                "Long": 25
            },
            "name": {
                "String": "alice"
            },
            "userId": {
                "String": "123456789012"
            }
        },
        "Parents": [
            {
                "EntityType": "PhotoApp::UserGroup",
                "EntityId": "alice_friends"
            },
            {
                "EntityType": "PhotoApp::UserGroup",
                "EntityId": "AVTeam"
            }
        ]
    },
    {
        "Identifier": {
            "EntityType": "PhotoApp::Photo",
            "EntityId": "vacationPhoto.jpg"
        },
        "Attributes": {
            "private": {
                "Boolean": false
            },
            "account": {
                "EntityIdentifier": {
                    "EntityType": "PhotoApp::Account",
                    "EntityId": "ahmad"
                }
            }
        },
        "Parents": []
```

```
    },
    {
        "Identifier": {
            "EntityType": "PhotoApp::UserGroup",
            "EntityId": "alice_friends"
        },
        "Parents": []
    },
    {
        "Identifier": {
            "EntityType": "PhotoApp::UserGroup",
            "EntityId": "AVTeam"
        },
        "Parents": []
    }
]
```

Length and size limits

Verified Permissions supports storage in the form of policy stores to hold your schema, policies, and policy templates. That storage causes Verified Permissions to impose some length and size limits that aren't relevant to Cedar.

Object	Verified Permissions limit (in bytes)	Cedar limit
Policy size[1]	10,000	None
Inline policy description	150	Not applicable to Cedar
Policy template size	10,000	None
Schema size	100,000	None
Entity type	200	None
Policy ID	64	None
Policy template ID	64	None

Object	Verified Permissions limit (in bytes)	Cedar limit
Entity ID	200	None
Policy store ID	64	Not applicable to Cedar

[1] There is a limit for policies per policy store in Verified Permissions based on the combined size of principals, actions, and resources of policies created in the policy store. The total size of all policies pertaining to a single resource can't exceed 200,000 bytes. For template-linked policies, the size of the policy template is counted only once, plus the size of each set of parameters used to instantiate each template-linked policy.

Amazon Verified Permissions upgrade to Cedar 4 FAQ

Amazon Verified Permissions is upgrading the version of Cedar it uses from version 2 to version 4. Cedar is the open-source language you use to write the policies, policy templates and schemas in your policy stores. With Cedar 4 support in Verified Permissions, you can use new features such as the `is` operator and entity tags to write more expressive policies.

Amazon Verified Permissions is automatically upgrading policy stores to Cedar 4. However, some policies, schemas and authorization requests written for Cedar 2 are incompatible with Cedar 4. If this is the case for your policy store, then we will not upgrade it automatically. You may need to make changes to your policies, policy templates, schemas or application code before you can upgrade to Cedar 4.

Topics

Why do I receive an error saying Cedar 2 is no longer supported when I call Amazon Verified Permissions?

Beginning April 2026, authorization APIs (`IsAuthorized`, `BatchIsAuthorized`, `IsAuthorizedWithToken` and `BatchIsAuthorizedWithToken`) are being disabled for policy stores which use Cedar 2. In this case, the following error message is returned:

```
Your policy store uses Cedar 2.x, which is no longer supported. See https://
docs.aws.amazon.com/verifiedpermissions/latest/userguide/cedar4-faq.html or contact AWS
 Support for more information.
```

If you encounter this error message when calling Verified Permissions, it means that your policy store is incompatible with Cedar 4, and so was unable to be migrated to the new version. To continue using Amazon Verified Permissions, create a new policy store, which will use Cedar 4, or contact Support.

Why are some policies, policy templates and schemas not compatible with Cedar 4?

The Cedar team has made several backwards-incompatible changes since Cedar 2, to fix bugs and simplify the language. These changes include:

- syntax changes for policies, policy templates and schemas
- a more precise policy validator, which detects more errors
- changes to the behaviour of built-in functions like `isInRange`

For a full list of backwards-incompatible changes, look for items marked with `(*)` in the Cedar changelog.

How do I tell whether my policy store is using Cedar 2 or Cedar 4?

You can check the version of Cedar your policy store uses using the Amazon Verified Permissions console, or using the GetPolicyStore operation.

> **Note**
>
> All policy stores in the same AWS account and region use the same version of Cedar.

Console

To check the Cedar version of a policy store (console)

1. Sign in to the AWS Management Console and open the Amazon Verified Permissions console at https://console.aws.amazon.com/verifiedpermissions/.

2. From the navigation pane, choose **Policy stores** and then choose the policy store that you want to check.
3. Choose **Settings** in the navigation pane.
4. In the **Details** box, locate the **Cedar version** field.

The field reads CEDAR_2 if your policy store is using Cedar 2, and CEDAR_4 if it uses Cedar 4.

CLI

To check the Cedar version of a policy store (AWS CLI)

1. Install and configure the AWS Command Line Interface (AWS CLI), if you haven't already. For information, see Installing or updating the latest version of the AWS CLI.
2. Use the `get-policy-store` command. In the following example, replace *`policy-store-id`* with the identifier of your policy store:

```
aws verifiedpermissions get-policy-store \
  --policy-store-id policy-store-id
```

The `cedarVersion` field in the output shows which version of Cedar the policy store is using. For example:

```
{
    "policyStoreId": "ABCDEFG12345678abcdefg",
    "arn": "arn:aws:verifiedpermissions::111122223333:policy-store/
ABCDEFG12345678abcdefg",
    "validationSettings": {
        "mode": "STRICT"
    },
    "createdDate": "2025-06-03T13:09:47.752255+00:00",
    "lastUpdatedDate": "2025-06-03T13:09:47.752255+00:00",
    "deletionProtection": "ENABLED",
    "cedarVersion": "CEDAR_2"
}
```

The field reads CEDAR_2 if your policy store is using Cedar 2, and CEDAR_4 if it uses Cedar 4.

How do I upgrade to Cedar 4?

Amazon Verified Permissions has already upgraded most customers to Cedar 4. If you have never created a policy store, then any new policy stores you create will use Cedar 4. If you are an existing customer, then we have likely already upgraded you to Cedar 4. See How do I tell whether my policy store is using Cedar 2 or Cedar 4? to check which version of Cedar your policy stores use.

If you have not been upgraded, then Verified Permissions detected a policy, policy template, schema or authorization request in one of your policy stores which is incompatible with Cedar 4. We will send you an email notification describing which resources are incompatible later in 2025. To upgrade sooner, open a case with Support.

> **Important**
>
> All policy stores in the same AWS account use the same version of Cedar. If one policy store in your account is incompatible with Cedar 4, then you can't use Cedar 4 in any policy store in that account.

Can I downgrade my policy store from Cedar 4 to Cedar 2?

No. If you experience issues after your policy store is upgraded to Cedar 4, open a case with Support.

Why am I receiving an error message saying my policy store is configured for Cedar 2?

Some features of Amazon Verified Permissions rely on the new features in Cedar 4. If your policy store does not use Cedar 4, then you can't use the following API fields:

- In the IsAuthorized, BatchIsAuthorized, IsAuthorizedWithToken and BatchIsAuthorizedWithToken operations:
  - `datetime`, `decimal` or `duration` values in the `attributes` or `context` fields

You can't use syntax or data types in policies, policy templates or schemas introduced after Cedar 2 until your policy store is upgraded.

How do I make my schema compatible with Cedar 4?

The Verified Permissions console can automatically fix some compatibility problems in your schema. If your schema cannot be automatically fixed, the console will show a list of errors for you to fix manually.

> **Important**
>
> The code editor in the Amazon Verified Permissions console always shows errors and warnings from Cedar 4, even if your policy store uses Cedar 2. You can continue to make schema updates that are not compatible with Cedar 4 using the **Save changes** button, or the Verified Permissions API.

To fix a schema using the console

1. Sign in to the AWS Management Console and open the Amazon Verified Permissions console at verifiedpermissions.
2. From the navigation pane, choose **Policy stores** and then choose the policy store you want to check.
3. Choose **Schema** in the navigation pane.
4. If your schema can be fixed automatically, you will see a banner reading “Click ‘Fix’ to preview a compatible version’. Select **Fix**.
5. Review the changes made to your schema, and click **Preview updated schema**.
6. Review the updated schema, and click **Save changes**.

If your schema can’t be fixed automatically, you can see a list of errors to fix yourself in the console.

1. Open the **Edit schema** page as described above.
2. Select **JSON mode**.
3. Hover over the red error icon in the gutter on the left-hand side of the code editor. The error message is displayed in a tooltip.

Here are some common errors you may encounter and how to resolve them:

failed to parse schema from JSON: `*field-name*`

With Cedar 2, you can include arbitrary fields in parts of schemas like type definitions, even if they do not have any meaning as part of a Cedar schema. In Cedar 4, this is no longer permitted. To resolve this error, remove the field called *field-name* from your JSON schema. For a list of valid schema fields, see the Cedar documentation.

unknown extension type `*extension-name*`

In Cedar 2, when you declare an attribute whose `type` is `Extension`, you can specify any value for the `name` field, whether or not the value is a valid extension type name. This is now an error with Cedar 4. To resolve it, replace *extension-name* with a valid extension type name. You can find a list of valid extension type names in the Cedar documentation.

If you are still unsure how to resolve the errors in your schema, contact Support

How do I make my policies and templates compatible with Cedar 4?

The Verified Permissions console shows you any errors in your policy or template which make it incompatible with Cedar 4.

To view a policy or template's errors in the console

1. Sign in to the AWS Management Console and open the Amazon Verified Permissions console at verifiedpermissions.
2. From the navigation pane, choose **Policy stores** and then choose the policy store you want to check.
3. Choose **Policies** or **Policy templates** in navigation pane, as appropriate.
4. Select the incompatible policy or template.
5. Select **Edit**
6. Hover over the red error icon in the gutter on the left-hand side of the code editor. The error message is displayed in a tooltip.

Here are some common errors you may encounter and how to resolve them:

empty set literals are forbidden in policies

In Cedar 2, you can use the syntax `mySet == []` to check whether a set is empty. With Cedar 4, policies using this syntax no longer validate against a schema. Replace `mySet == []` in your policy with `mySet.isEmpty()`.

Document history for the Amazon Verified Permissions User Guide

The following table describes the documentation releases for Verified Permissions.

Change	Description	Date
Hard delete for policy store aliases	You can now hard delete policy store aliases for immediate deletion, bypassing the 24-hour `PendingDeletion` reservation period.	April 17, 2026
Policy names and policy template names	You can now assign names to policies and policy templates to reference them by friendly names.	March 4, 2025
Policy store aliases	You can now create policy store aliases to refer to your policy stores by friendly names.	February 26, 2025
New AWS managed policies	You can now use the `AmazonVerifiedPermissionsFullAccess` and `AmazonVerifiedPermissionsReadOnlyAccess` IAM managed policies with Verified Permissions.	October 11, 2024
OIDC identity sources	You can now authorize users from OpenID Connect (OIDC) identity providers.	June 8, 2024

Batch authorization with identity source tokens	You can now authorize users from a Amazon Cognito user pool in a single `BatchIsAuthorizedWithToken` API request.	April 5, 2024
Creating a policy store with API Gateway	You can now create a policy store from an existing API and Amazon Cognito user pool.	April 1, 2024
Context concepts and example	Added information about context in authorization requests with Verified Permissions.	February 1, 2024
Authorization concepts and example	Added information about authorization requests with Verified Permissions.	February 1, 2024
AWS CloudFormation integration	Verified Permissions supports creating identity sources, policies, policy stores, and policy templates in CloudFormation.	June 30, 2023
Initial release	Initial release of the Amazon Verified Permissions User Guide	June 13, 2023