

User Guide

Amazon Verified Permissions



Copyright © 2025 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon Verified Permissions: User Guide

Copyright © 2025 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

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

Table of Contents

What is Amazon Verified Permissions?	, 1
Authorization in Verified Permissions	. 1
Cedar policy language	. 2
Benefits of Verified Permissions	. 2
Accelerate application development	. 2
More secure applications	. 2
End-user features	. 3
Related services	. 3
Accessing Verified Permissions	. 3
Pricing for Verified Permissions	5
Getting started with policy stores	6
Prerequisites	7
Step 1: Create a PhotoFlash policy store	. 9
Step 2: Create a policy	. 9
Step 3: Testing a policy store	10
Step 4: Clean up resources	12
Designing an authorization model	13
No single correct model	14
Returning errors	14
Focus on resources	15
Consider multi-tenancy	16
Comparing shared policy stores and per-tenant policy stores	18
How to choose	19
Policy stores	20
Creating policy stores	20
Creating a policy store using Rust	28
API-linked policy stores	33
How it works	35
Considerations	37
Adding ABAC	38
Moving to production	39
Troubleshooting	42
Deleting policy stores	44
Policy store schema	47

Editing schema	49
Policy validation mode	52
Policies	54
Creating static policies	. 55
Editing static policies	57
	59
Evaluate example context	. 61
Testing policies	. 67
Example policies	. 69
Uses bracket notation to reference token attributes	. 70
Uses dot notation to reference attributes	. 70
Reflects Amazon Cognito ID token attributes	. 71
Reflects OIDC ID token attributes	. 71
Reflects Amazon Cognito access token attributes	. 72
Reflects OIDC access token attributes	. 72
Policy templates and template-linked policies	. 73
Creating policy templates	. 73
Creating template-linked policies	. 75
Editing policy templates	. 77
Example template-linked policies	. 79
PhotoFlash examples	. 79
DigitalPetStore examples	. 80
TinyToDo examples	81
Identity sources	82
Choosing the right identity provider	. 83
Working with Amazon Cognito identity sources	. 83
Creating identity sources	. 86
Editing identity sources	. 89
Mapping tokens to schema	91
Client and audience validation	102
Working with OIDC identity sources	105
Creating identity sources	106
Editing identity sources	109
Mapping tokens to schema	111
Client and audience validation	118
Integrations	122

Using Express	122
Prerequisites	123
Setting up the integration	123
Configuring authorization	124
Implementing the authorization middleware	127
Testing the integration	127
Troubleshooting	128
Next steps	128
Authorize requests	129
API operations	130
Test model	131
Integrating with applications	133
Security	136
Data protection	136
Data encryption	138
Identity and access management	138
Audience	139
Authenticating with identities	139
Managing access using policies	142
How Amazon Verified Permissions works with IAM	145
IAM policies for Verified Permissions	151
Identity-based policy examples	154
AWS managed policies	157
Troubleshooting	160
Compliance validation	162
Resilience	163
Monitoring	165
CloudTrail logs	165
Verified Permissions information in CloudTrail	165
Understanding Verified Permissions log file entries	166
Working with AWS CloudFormation	184
Verified Permissions and AWS CloudFormation templates	184
AWS CDK constructs	185
Learn more about AWS CloudFormation	185
Using AWS PrivateLink	186
Considerations	186

Create an interface endpoint	
Create an endpoint policy	187
Quotas	
Quotas for resources	
Template-linked policy size example	190
Quotas for hierarchies	192
Quotas for operations per second	193
Terms & concepts	198
Authorization model	199
Authorization request	199
Authorization response	199
Considered policies	199
Context data	199
Determining policies	200
Entity data	200
Permissions, authorization, and principals	200
Policy enforcement	200
Policy store	200
Satisfied policies	201
Differences with Cedar	201
Namespace definition	201
Policy template support	201
Schema support	202
Action groups definition	202
Entity formatting	202
Length and size limits	207
Cedar v4 FAQ	209
What is the current state on the upgrade?	209
Do I need to do anything right now?	209
Does the upgrade of the console impact the authorization service?	209
What are the breaking changes in Cedar v3 and Cedar v4?	210
When will the upgrade to Cedar v4 be complete?	210
Document history	211

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 <u>Implementing a PDP by using Amazon Verified Permissions</u> in AWS Prescriptive Guidance.

Topics

- Authorization in Verified Permissions
- <u>Cedar policy language</u>
- <u>Benefits of Verified Permissions</u>
- <u>Related services</u>
- <u>Accessing Verified Permissions</u>
- Pricing for Verified Permissions

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 AWS 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 2.4.

For more information about Cedar, see the following:

- <u>Cedar policy language Reference Guide</u>
- 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 <u>the</u> 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 <u>Amazon Cognito</u> <u>Developer Guide</u>.
- 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 <u>API Gateway Developer Guide</u>.
- 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 <u>AWS IAM Identity Center</u> <u>User Guide</u>.

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 <u>How to</u> <u>sign in to AWS</u> 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 <u>AWS Command Line Interface</u> (AWS CLI) and the <u>AWS Tools for Windows PowerShell</u>. For information about installing and using the AWS CLI, see the <u>AWS Command Line Interface User Guide</u>. For information about installing and using the Tools for Windows PowerShell, see the <u>AWS Tools for PowerShell User Guide</u>.

- 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
- <u>AWS SDK for JavaScript</u>
- 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 AWS CloudFormation. Constructs, or reusable cloud components, can be used to create AWS 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 <u>Amazon Verified Permissions</u> pricing.

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


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 <u>policy stores</u>, 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

If you do not have an AWS account, complete the following steps to create one.

To sign up for an AWS account

- 1. Open https://portal.aws.amazon.com/billing/signup.
- 2. Follow the online instructions.

Part of the sign-up procedure involves receiving a phone call or text message and entering a verification code on the phone keypad.

When you sign up for an AWS account, an AWS account root user is created. The root user has access to all AWS services and resources in the account. As a security best practice, assign

administrative access to a user, and use only the root user to perform <u>tasks that require root</u> user access.

AWS sends you a confirmation email after the sign-up process is complete. At any time, you can view your current account activity and manage your account by going to <u>https://aws.amazon.com/</u> and choosing **My Account**.

Create a user with administrative access

After you sign up for an AWS account, secure your AWS account root user, enable AWS IAM Identity Center, and create an administrative user so that you don't use the root user for everyday tasks.

Secure your AWS account root user

1. Sign in to the <u>AWS Management Console</u> as the account owner by choosing **Root user** and entering your AWS account email address. On the next page, enter your password.

For help signing in by using root user, see <u>Signing in as the root user</u> in the AWS Sign-In User Guide.

2. Turn on multi-factor authentication (MFA) for your root user.

For instructions, see Enable a virtual MFA device for your AWS account root user (console) in the IAM User Guide.

Create a user with administrative access

1. Enable IAM Identity Center.

For instructions, see <u>Enabling AWS IAM Identity Center</u> in the AWS IAM Identity Center User *Guide*.

2. In IAM Identity Center, grant administrative access to a user.

For a tutorial about using the IAM Identity Center directory as your identity source, see <u>Configure user access with the default IAM Identity Center directory</u> in the AWS IAM Identity Center User Guide. • To sign in with your IAM Identity Center user, use the sign-in URL that was sent to your email address when you created the IAM Identity Center user.

For help signing in using an IAM Identity Center user, see <u>Signing in to the AWS access portal</u> in the AWS Sign-In User Guide.

Assign access to additional users

1. In IAM Identity Center, create a permission set that follows the best practice of applying leastprivilege permissions.

For instructions, see <u>Create a permission set</u> in the AWS IAM Identity Center User Guide.

2. Assign users to a group, and then assign single sign-on access to the group.

For instructions, see <u>Add groups</u> in the AWS IAM Identity Center User 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 <u>Verified Permissions console</u>, 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**.
- For Resources scope, choose Specific resource, then for Specify entity type, choose PhotoFlash::Album, and for Specify entity identifier, enter alice-favorites-album.
- 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 <u>authorization</u> 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	Authoriza tion decision
PhotoFlas h::User bob	PhotoFlas h::Account alice-account	PhotoFlas h::Album alice-fav orites-album	N/A	PhotoFlas h::Action ::"Update Album"	Deny
PhotoFlas h::User alice	PhotoFlas h::Account alice-account	PhotoFlas h::Photo photo.jpeg	PhotoFlas h::Account bob-account	PhotoFlas h::Action ::"ViewPh oto"	Deny
PhotoFlas h::User alice	PhotoFlas h::Account alice-account	PhotoFlas h::Photo photo.jpeg	PhotoFlas h::Account alice-account	PhotoFlas h::Action ::"ViewPh oto"	Allow

Principal value	Principal Account: Entity value	Resource value	Resource parent value	Action	Authoriza tion decision
PhotoFlas h::User alice	PhotoFlas h::Account alice-account	PhotoFlas h::Photo bob-photo .jpeg	PhotoFlas h::Album Bob-Vacat ion-Album	PhotoFlas h::Action ::"Delete Photo"	Deny

Step 4: Clean up resources

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

To delete a policy store

- 1. In the <u>Verified Permissions console</u>, 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 <u>best practices</u> 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
- <u>Multi-tenancy considerations</u>

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 filesharing 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 <u>getting started</u> 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
listGroupMembershi psForUser	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 <u>Amazon Verified Permissions multi-tenant design considerations</u> 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 <u>policy templates</u> and a <u>policy store</u> 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 multitenancy 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 pertenant policy store models.

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

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 one namespace 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 <u>Namespaces</u> in the *Cedar policy language Reference Guide*.

Topics

- <u>Creating Verified Permissions policy stores</u>
- API-linked policy stores
- Deleting policy stores

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 <u>Verified Permissions console</u>, 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 <u>Namespace definition</u>.
- 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 <u>Authorization with Amazon Verified</u> <u>Permissions</u> 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*.

i 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 <u>authorizer</u> 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 <u>Setting up a stage for a REST API in</u> 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 <u>Mapping Amazon Cognito tokens to</u> 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.
 - 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 <u>Use API Gateway Lambda authorizers</u> 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 <u>User pool authentication flow</u> 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 <u>Amazon Verified Permissions example template-linked policies</u>.

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 <u>Amazon Verified</u> <u>Permissions policy store schema</u>.

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

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/
PSEXAMPLEabcdefg11111",
    "createdDate": "2023-05-16T17:41:29.103459+00:00",
    "lastUpdatedDate": "2023-05-16T17:41:29.103459+00:00",
    "policyStoreId": "PSEXAMPLEabcdefg11111"
}
```

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 <u>aws-sdk-verifiedpermissions</u> 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 <u>AWS CLI</u> 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 <u>Configuring settings for the AWS CLI</u> and <u>Configuration and credential file settings in the AWS CLI</u>.
- 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 <u>authorization request</u> 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_stc
    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.

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 <u>Amazon Cognito</u>, 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 <u>Cedar policy language</u>. 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 <u>Amazon Verified Permissions - Quick Start Overview and Demo</u> 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 <u>AWS Cloud Development Kit (CDK)</u>. For more information, see <u>Moving to production with AWS CloudFormation</u>.


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 <u>IsAuthorizedWithToken</u> 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 <u>Amazon Cognito</u> 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 <u>API Gateway</u> REST API. Verified Permissions defines actions from API paths and API methods, for example MyAPI::Action::get /photo.
- A Lambda function and a <u>Lambda authorizer</u> 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 <u>authorizer</u>.



- 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 <u>Using tokens with user pools</u> 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 <u>IsAuthorizedWithToken</u> 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 <u>Authorization with Amazon Verified Permissions</u> 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 <u>namespace</u> 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 AWS 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 <u>Moving to production with AWS CloudFormation</u> 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 <u>Using tokens with user pools</u> 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 <u>Mapping Amazon Cognito tokens to schema</u> and <u>Mapping OIDC tokens to schema</u> 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:

- <u>AWS::VerifiedPermissions::PolicyStore</u>: Copy your schema to the SchemaDefinition object. Escape " characters as \".
- 2. <u>AWS::VerifiedPermissions::IdentitySource</u>: Copy values from the output of <u>GetIdentitySource</u> from your test policy store and modify as needed.
- 3. One or more of <u>AWS::VerifiedPermissions::Policy</u>: Copy your policy statement to the Definition object. Escape " characters as \".
- 4. <u>AWS::Lambda::Function</u>, <u>AWS::IAM::Role</u>, <u>AWS::IAM::Policy</u>, <u>AWS::ApiGateway::Authorizer</u>, 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::
\space{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"
            1
        },
        "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
- I attached the Lambda authorizer to my API but it's not generating authorization requests
- I received an unexpected authorization decision and want to review the authorization logic
- I want to find logs from my Lambda authorizer
- <u>My Lambda authorizer doesn't exist</u>
- My API is in a private VPC and can't invoke the authorizer
- I want to process additional user attributes in my authorization model
- I want to add new actions, action context attributes, or resource attributes

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 <u>Enabling API caching to enhance responsiveness</u> 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 <u>Deploying a REST API</u> 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, <u>locate your function</u> 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 <u>Using Amazon CloudWatch Logs with AWS</u> <u>Lambda</u> 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 AWS 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 <u>Troubleshooting CloudFormation</u> 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 <u>Mapping Amazon</u> Cognito tokens to schema and <u>Mapping OIDC tokens to schema</u>.

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 <u>start again with a new policy store</u> or you can update your existing policy store.

To update your existing policy store, <u>locate your function</u>. Examine the logic in the automaticallygenerated function and update it to process the new actions, attributes, or context. Then <u>edit your</u> <u>schema</u> 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 in the policy store. 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 <u>Verified Permissions console</u>. 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 PSEXAMPLEabcdefg11111
```

If successful, this command produces no output.

🚯 Note

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

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

Amazon Verified Permissions policy store schema

A <u>schema</u> 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 <u>Cedar schema format</u> in the Cedar policy language Reference Guide.

1 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 <u>PutSchema</u> 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 <u>Verified Permissions console</u>. 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.
- (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.
- (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 <u>Verified Permissions console</u>. 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 <u>schema</u> 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 <u>UpdatePolicyStore</u> operation and specifying a different value for the <u>ValidationSettings</u> parameter.

```
$ aws verifiedpermissions update-policy-store \
    --validation-settings "mode=OFF",
    --policy-store-id PSEXAMPLEabcdefg11111
{
        "createdDate": "2023-05-17T18:36:10.134448+00:00",
        "lastUpdatedDate": "2023-05-17T18:36:10.134448+00:00",
        "policyStoreId": "PSEXAMPLEabcdefg11111",
        "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 <u>Cedar policy validation against schema</u> in the Cedar policy language Reference Guide.

A 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
- Editing Amazon Verified Permissions static policies
- Adding context
- Using the Amazon Verified Permissions test bench

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.
- 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 <u>Basic policy construction in Cedar</u> 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 <u>CreatePolicy</u> 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 PSEXAMPLEabcdefg11111
{
    "Arn": "arn:aws:verifiedpermissions::123456789012:policy/PSEXAMPLEabcdefg11111/
SPEXAMPLEabcdefg11111",
    "createdDate": "2023-05-16T20:33:01.730817+00:00",
    "lastUpdatedDate": "2023-05-16T20:33:01.730817+00:00",
    "policyId": "SPEXAMPLEabcdefg11111",
    "policyStoreId": "PSEXAMPLEabcdefg11111",
```

}

"policyType": "STATIC"

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 <u>policy validation</u> 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 <u>UpdatePolicy</u> 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 PSEXAMPLEabcdefg11111
{
        "createdDate": "2023-06-12T20:33:37.382907+00:00",
        "lastUpdatedDate": "2023-06-12T20:33:37.382907+00:00",
        "policyId": "SPEXAMPLEabcdefg11111",
        "policyStoreId": "PSEXAMPLEabcdefg11111",
        "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 <u>sample policy store</u>. 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 <u>Mapping Amazon Cognito access tokens</u> and <u>Mapping OIDC access tokens</u>.

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 <u>commonType</u> 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(11112223333) &&
    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*"</pre>
```

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 <u>IsAuthorized</u>. 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": {
        "entityMap": {
        "entity
```

```
"AccountCodes": {
      "set": [
        {"long": 111122223333},
        {"long": 444455556666},
        {"long": 123456789012}
      1
    },
    "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 <u>authorization requests</u> 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 <u>Basic policy construction in</u> Cedar in the Cedar policy language Reference Guide.

1 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 <u>Amazon Verified Permissions</u> - <u>Policy Creation and Testing (Primer Series #3)</u> 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 <u>Verified Permissions console</u>. 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.

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.

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 <u>Basic policy construction in</u> <u>Cedar</u> 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
- <u>Reflects Amazon Cognito ID token attributes</u>
- <u>Reflects OIDC ID token attributes</u>
- <u>Reflects Amazon Cognito access token attributes</u>
- <u>Reflects OIDC access token attributes</u>

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 <u>Mapping</u> 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 <u>Mapping</u> 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 <u>Mapping</u> 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 <u>Mapping</u> 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 <u>Mapping</u> 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 <u>Mapping</u> 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 <u>Cedar policy templates</u> in the Cedar policy language Reference Guide.

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 <u>Amazon Verified Permissions example template-</u> <u>linked policies</u>.

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 <u>CreatePolicyTemplate</u> 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"
```

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 <u>Amazon Verified Permissions example template-linked</u> 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 PSEXAMPLEabcdefg11111
{
        "createdDate": "2023-05-22T18:57:53.298278+00:00",
        "lastUpdatedDate": "2023-05-22T18:57:53.298278+00:00",
        "policyId": "TPEXAMPLEabcdefg11111",
        "policyStoreId": "PSEXAMPLEabcdefg11111",
        "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 <u>UpdatePolicy</u> 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": "PSEXAMPLEabcdefg11111",
    "policyTemplateId": "PTEXAMPLEabcdefg11111"
}
```

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-1581886866f31|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|5ae@c4b1-6de8-4dff-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 <u>OpenID Connect (OIDC)</u> 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 <u>context</u>. 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).

i 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 <u>Authorize API Gateway APIs</u> <u>using Amazon Verified Permissions with Amazon Cognito or bring your own identity provider</u> on the *AWS Security Blog*.

Topics

- <u>Choosing the right identity provider</u>
- 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-EXAMPLE2222 as us-west-2_EXAMPLE|a1b2c3d4-5678-90ab-cdef-EXAMPLE2222

User pool token *claims* can contain attributes, scopes, groups, client IDs, and custom data. <u>Amazon Cognito JWTs</u> 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.

<u> Important</u>

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 <u>the JSON Web Token RFC 7519</u> 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 <u>Ending user sessions</u> 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.

1 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 <u>IsAuthorizedWithToken</u> will fail with a ValidationException.

For more information about mapping claims, see <u>Mapping Amazon Cognito tokens to schema</u>. For more information about authorization for Amazon Cognito users, see <u>Authorization with Amazon</u> <u>Verified Permissions</u> 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
- <u>Client and audience validation for Amazon Cognito</u>

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 <u>create a new policy store</u> 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 <u>create-identity-</u> <u>source</u> in the AWS CLI or <u>CreateIdentitySource</u> 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 <u>API-linked policy store</u> 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 <u>Creating Amazon Verified Permissions static policies</u>.

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 <u>IsAuthorizedWithToken</u> 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 <u>CreateIdentitySource</u> operation. The following example creates an identity source that can access authenticated identities from a Amazon Cognito user pool.

 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": "ISEXAMPLEabcdefg11111",
    "lastUpdatedDate": "2023-05-19T20:30:28.214829+00:00",
    "policyStoreId": "PSEXAMPLEabcdefg11111"
}
```

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 <u>Mapping Amazon Cognito tokens to schema</u>.

🚯 Note

When you create an <u>API-linked policy store</u> 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 <u>Creating Amazon Verified Permissions static policies</u>.

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 <u>IsAuthorizedWithToken</u> 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 <u>Authorization with Amazon Verified Permissions</u> 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 <u>UpdateIdentitySource</u> operation. The following example updates the specified identity source to use a different Amazon Cognito user pool.

 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": "ISEXAMPLEabcdefg11111",
    "lastUpdatedDate": "2023-05-19T20:30:28.214829+00:00",
    "policyStoreId": "PSEXAMPLEabcdefg11111"
}
```

🚯 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 <u>Guided setup</u> 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

<u>API-linked policy stores</u> and policy stores with an identity source that were created through <u>Guided</u> <u>setup</u> 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 <u>IsAuthorizedWithToken</u> or <u>BatchIsAuthorizedWithToken</u> 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 <u>Authorization with Amazon Verified Permissions</u> in the *Amazon Cognito Developer Guide*.

Topics

- <u>Mapping ID tokens to schema</u>
- <u>Mapping access tokens</u>
- Alternative notation for Amazon Cognito colon-delimited claims
- <u>Things to know about schema mapping</u>

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 <u>OIDC relying-party libraries</u>. 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 <u>Using the API and endpoints</u> 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 <u>Standard</u> 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 <u>Pre token</u> <u>generation Lambda trigger</u>. 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 <u>Using</u> 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 <u>Editing policy</u>

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 <u>Reflects Amazon Cognito ID token</u> <u>attributes</u>.

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 <u>Pre token</u> <u>generation Lambda trigger</u>. 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-T0KEN111111",
    "event_id": "bda909cb-3e29-4bb8-83e3-ce6808f49011",
```

```
User Guide
```

```
"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 <u>Editing policy</u> <u>store schemas</u>.

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

For an example policy that will validate against this schema, see <u>Reflects Amazon Cognito access</u> 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 <u>bracket notation</u>. 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 <u>Uses dot</u> <u>notation to reference attributes</u>.

Things to know about schema mapping

Attribute mapping differs between token types

In access token authorization, Verified Permissions maps claims to <u>context</u>. 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 groupmembership 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 <u>commonTypes</u> attributes. For more information, see <u>Mapping access tokens</u>.

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 <u>API-linked policy store</u>, 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 <u>Uses bracket notation to reference</u> <u>token attributes</u>.

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 <u>app client</u> 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 <u>aws-jwt-verify</u> 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/openidconfiguration 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-EXAMPLE1111".

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

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 <u>create a new policy store</u> 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 <u>create-identity-</u> <u>source</u> in the AWS CLI or <u>CreateIdentitySource</u> 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 <u>parent</u> 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 <u>Creating Amazon Verified Permissions static policies</u>.

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 <u>IsAuthorizedWithToken</u> 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 <u>CreateIdentitySource</u> operation. The following example creates an identity source that can access authenticated identities from a an OIDC identity provider(IdP).

 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 <u>IsAuthorizedWithToken</u> 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 <u>Verified Permissions console</u>. 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 <u>UpdateIdentitySource</u> operation. The following example updates the specified identity source to use a different OIDC provider.

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": "PSEXAMPLEabcdefg11111"
```

🚯 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 <u>Guided setup</u> 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

<u>API-linked policy stores</u> and policy stores with an identity source that were created through <u>Guided</u> <u>setup</u> 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 <u>IsAuthorizedWithToken</u> or <u>BatchIsAuthorizedWithToken</u> 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
- Mapping access tokens
- Things to know about schema mapping

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 <u>standard OIDC attributes</u>, 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": {
```

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-T0KEN2222222",
    "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 <u>Editing policy</u> <u>store schemas</u>.

```
{
   "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 <u>Reflects OIDC access token</u> <u>attributes</u>.

Things to know about schema mapping

Attribute mapping differs between token types

In access token authorization, Verified Permissions maps claims to <u>context</u>. 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 groupmembership 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 <u>commonTypes</u> attributes. For more information, see <u>Mapping access tokens</u>.

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"]

1 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 <u>Uses bracket notation to reference</u> <u>token attributes</u>.

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 <u>aws-jwt-verify</u> 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-expressis 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
- <u>Node.js</u> and <u>npm</u> 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 <u>namespace</u> 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 <u>OpenAPI</u> <u>specifications</u> 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 <u>express-openapi-generator</u> 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 <u>verifiedpermissions/examples</u> 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 <u>Working with AWS</u> <u>CloudFormation</u>.

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 <u>create-policy</u> in the AWS CLI reference. If you need help formatting the policies, there's a script called convert_cedar_policies.sh in the <u>verifiedpermissions/examples</u> 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-storeid 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 <u>authorization-for-expressis</u> 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 <u>sample policy store</u>. 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 <u>Use Amazon Verified Permissions for fine-grained authorization at scale</u> 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 validates the entities in your request against your policy store schema. Verified Permissions then 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 <u>quota usage</u> 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 <u>quota usage</u> 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 <u>Using the Amazon</u> <u>Verified Permissions test bench</u> 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 <u>IsAuthorized</u>, <u>IsAuthorizedWithToken</u>, and <u>BatchIsAuthorized</u> as <u>signed AWS API</u> <u>requests</u> to Regional <u>service endpoints</u>. 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 <u>Verified Permissions console</u>. Create a policy store from the **Sample policy store** with the name **DigitalPetStore**.
- 2. Select **Test bench** in your new policy store.
- Populate your test bench request from <u>IsAuthorized</u> 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 <u>Verified Permissions console</u>. Create a policy store from the **Sample policy store** with the name **DigitalPetStore**.
- 2. Note the **Policy store ID** of your new policy store.
- From <u>IsAuthorized</u> 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 <u>example</u>.
- 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

- 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 <u>identity source</u> linked to your policy store or from another <u>OpenID</u>
 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 <u>Tools for Amazon Web Services</u>.

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
- <u>AWS SDK for JavaScript</u>
- 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 <u>Simplify fine-</u> 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

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 <u>shared responsibility model</u> 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 <u>AWS</u>
 <u>Compliance Programs</u>. 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
- Identity and access management for Amazon Verified Permissions
- <u>Compliance validation for Amazon Verified Permissions</u>
- <u>Resilience in Amazon Verified Permissions</u>

Data protection in Amazon Verified Permissions

The AWS <u>shared responsibility model</u> 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 <u>Data Privacy FAQ</u>. For information about data protection in Europe, see the <u>AWS Shared Responsibility Model and GDPR</u> 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 <u>Federal Information Processing Standard (FIPS) 140-2</u>.
- 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, so the use of a customer managed key is neither necessary nor supported.

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

- <u>Audience</u>
- Authenticating with identities
- Managing access using policies
- How Amazon Verified Permissions works with IAM
- IAM policies for Verified Permissions
- Identity-based policy examples for Amazon Verified Permissions
- AWS managed policies for Amazon Verified Permissions
- <u>Troubleshooting Amazon Verified Permissions identity and access</u>

Audience

How you use AWS Identity and Access Management (IAM) differs, depending on the work that you do in Verified Permissions.

Service user – If you use the Verified Permissions service to do your job, then your administrator provides you with the credentials and permissions that you need. As you use more Verified Permissions features to do your work, you might need additional permissions. Understanding how access is managed can help you request the right permissions from your administrator. If you cannot access a feature in Verified Permissions, see <u>Troubleshooting Amazon Verified Permissions</u> identity and access.

Service administrator – If you're in charge of Verified Permissions resources at your company, you probably have full access to Verified Permissions. It's your job to determine which Verified Permissions features and resources your service users should access. You must then submit requests to your IAM administrator to change the permissions of your service users. Review the information on this page to understand the basic concepts of IAM. To learn more about how your company can use IAM with Verified Permissions, see <u>How Amazon Verified Permissions works with IAM</u>.

IAM administrator – If you're an IAM administrator, you might want to learn details about how you can write policies to manage access to Verified Permissions. To view example Verified Permissions identity-based policies that you can use in IAM, see <u>Identity-based policy examples for Amazon</u> <u>Verified Permissions</u>.

Authenticating with identities

Authentication is how you sign in to AWS using your identity credentials. You must be *authenticated* (signed in to AWS) as the AWS account root user, as an IAM user, or by assuming an IAM role.

You can sign in to AWS as a federated identity by using credentials provided through an identity source. AWS IAM Identity Center (IAM Identity Center) users, your company's single sign-on authentication, and your Google or Facebook credentials are examples of federated identities. When you sign in as a federated identity, your administrator previously set up identity federation using IAM roles. When you access AWS by using federation, you are indirectly assuming a role.

Depending on the type of user you are, you can sign in to the AWS Management Console or the AWS access portal. For more information about signing in to AWS, see <u>How to sign in to your AWS</u> <u>account</u> in the AWS Sign-In User Guide.

If you access AWS programmatically, AWS provides a software development kit (SDK) and a command line interface (CLI) to cryptographically sign your requests by using your credentials. If you don't use AWS tools, you must sign requests yourself. For more information about using the recommended method to sign requests yourself, see <u>AWS Signature Version 4 for API requests</u> in the *IAM User Guide*.

Regardless of the authentication method that you use, you might be required to provide additional security information. For example, AWS recommends that you use multi-factor authentication (MFA) to increase the security of your account. To learn more, see <u>Multi-factor authentication</u> in the AWS IAM Identity Center User Guide and <u>AWS Multi-factor authentication in IAM</u> in the IAM User Guide.

AWS account root user

When you create an AWS account, you begin with one sign-in identity that has complete access to all AWS services and resources in the account. This identity is called the AWS account *root user* and is accessed by signing in with the email address and password that you used to create the account. We strongly recommend that you don't use the root user for your everyday tasks. Safeguard your root user credentials and use them to perform the tasks that only the root user can perform. For the complete list of tasks that require you to sign in as the root user, see <u>Tasks that require root</u> user credentials in the *IAM User Guide*.

Federated identity

As a best practice, require human users, including users that require administrator access, to use federation with an identity provider to access AWS services by using temporary credentials.

A *federated identity* is a user from your enterprise user directory, a web identity provider, the AWS Directory Service, the Identity Center directory, or any user that accesses AWS services by using credentials provided through an identity source. When federated identities access AWS accounts, they assume roles, and the roles provide temporary credentials.

For centralized access management, we recommend that you use AWS IAM Identity Center. You can create users and groups in IAM Identity Center, or you can connect and synchronize to a set of users and groups in your own identity source for use across all your AWS accounts and applications. For information about IAM Identity Center, see <u>What is IAM Identity Center?</u> in the AWS IAM Identity Center User Guide.

specific use cases that require long-term credentials with IAM users, we recommend that you rotate access keys. For more information, see <u>Rotate access keys regularly for use cases that require long-</u>term credentials in the *IAM User Guide*.

An <u>IAM group</u> is an identity that specifies a collection of IAM users. You can't sign in as a group. You can use groups to specify permissions for multiple users at a time. Groups make permissions easier to manage for large sets of users. For example, you could have a group named *IAMAdmins* and give that group permissions to administer IAM resources.

Users are different from roles. A user is uniquely associated with one person or application, but a role is intended to be assumable by anyone who needs it. Users have permanent long-term credentials, but roles provide temporary credentials. To learn more, see <u>Use cases for IAM users</u> in the *IAM User Guide*.

IAM roles

An <u>IAM role</u> 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 <u>switching roles</u>. 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 <u>Using IAM roles</u> 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 <u>Create a role for a third-party identity provider</u> (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 <u>Permission sets</u> 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 <u>How IAM roles differ from resource-based policies</u> 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 <u>Use an IAM role to grant permissions to applications running on Amazon EC2 instances</u> in the *IAM User Guide*.

To learn whether to use IAM roles or IAM users, see <u>When to create an IAM role (instead of a user)</u> 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 is an object in AWS that, when associated with an identity or resource, defines their permissions. AWS evaluates these policies when a principal (user, root user, or role session) makes a request. Permissions in the policies determine whether the request is allowed or denied. Most policies are stored in AWS as JSON documents. For more information about the structure and contents of JSON policy documents, see <u>Overview of JSON policies</u> in the *IAM User Guide*.

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**.

By default, users and roles have no permissions. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

IAM policies define permissions for an action regardless of the method that you use to perform the operation. For example, suppose that you have a policy that allows the iam:GetRole action. A user with that policy can get role information from the AWS Management Console, the AWS CLI, or the AWS API.

Identity-based policies

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*.

Identity-based policies can be further categorized as *inline policies* or *managed policies*. Inline policies are embedded directly into a single user, group, or role. Managed policies are standalone policies that you can attach to multiple users, groups, and roles in your AWS account. Managed policies include AWS managed policies and customer managed policies. To learn how to choose between a managed policy or an inline policy, see <u>Choose between managed policies and inline policies</u> in the *IAM User Guide*.

Resource-based policies

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 <u>specify a principal</u> in a resource-based policy. Principals can include accounts, users, roles, federated users, or AWS services.

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 <u>Access control list (ACL) overview</u> in the *Amazon Simple Storage Service Developer Guide*.

Other policy types

AWS supports additional, less-common policy types. These policy types can set the maximum permissions granted to you by the more common policy types.

- Permissions boundaries A permissions boundary is an advanced feature in which you set the maximum permissions that an identity-based policy can grant to an IAM entity (IAM user or role). You can set a permissions boundary for an entity. The resulting permissions are the intersection of an entity's identity-based policies and its permissions boundaries. Resource-based policies that specify the user or role in the Principal field are not limited by the permissions boundary. An explicit deny in any of these policies overrides the allow. For more information about permissions boundaries, see <u>Permissions boundaries for IAM entities</u> in the *IAM User Guide*.
- Service control policies (SCPs) SCPs are JSON policies that specify the maximum permissions for an organization or organizational unit (OU) in AWS Organizations. AWS Organizations is a service for grouping and centrally managing multiple AWS accounts that your business owns. If you enable all features in an organization, then you can apply service control policies (SCPs) to any or all of your accounts. The SCP limits permissions for entities in member accounts, including each AWS account root user. For more information about Organizations and SCPs, see <u>Service</u> control policies in the AWS Organizations User Guide.
- Resource control policies (RCPs) RCPs are JSON policies that you can use to set the maximum available permissions for resources in your accounts without updating the IAM policies attached to each resource that you own. The RCP limits permissions for resources in member accounts and can impact the effective permissions for identities, including the AWS account root user, regardless of whether they belong to your organization. For more information about Organizations and RCPs, including a list of AWS services that support RCPs, see <u>Resource control policies (RCPs)</u> in the AWS Organizations User Guide.
- Session policies Session policies are advanced policies that you pass as a parameter when you programmatically create a temporary session for a role or federated user. The resulting session's permissions are the intersection of the user or role's identity-based policies and the session policies. Permissions can also come from a resource-based policy. An explicit deny in any of these policies overrides the allow. For more information, see <u>Session policies</u> 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 <u>AWS services that work with IAM</u> in the *IAM User Guide*.

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 <u>Define custom IAM permissions with customer managed policies</u> 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. 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 <u>IAM JSON policy elements reference</u> in the *IAM User Guide*.

Identity-based policy examples for Verified Permissions

To view examples of Verified Permissions identity-based policies, see <u>Identity-based policy</u> <u>examples for Amazon Verified Permissions</u>.

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 <u>specify a principal</u> 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. Adding a cross-account principal to a resource-based policy is only half of establishing the trust relationship. When the principal and the resource

are in different AWS accounts, an IAM administrator in the trusted account must also grant the principal entity (user or role) permission to access the resource. They grant permission by attaching an identity-based policy to the entity. However, if a resource-based policy grants access to a principal in the same account, no additional identity-based policy is required. For more information, see <u>Cross account resource access in IAM</u> 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. Policy actions usually have the same name as the associated AWS API operation. There are some exceptions, such as *permission-only actions* that don't have a matching API operation. There are also some operations that require multiple actions in a policy. These additional actions are called *dependent actions*.

Include actions in a policy to grant permissions to perform the associated operation.

To see a list of Verified Permissions actions, see <u>Actions defined by Amazon Verified Permissions</u> 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 <u>Identity-based policy</u> 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. Statements must include either a Resource or a NotResource element. As a best practice, specify a resource using its <u>Amazon Resource Name (ARN)</u>. You can do this for actions that support a specific resource type, known as *resource-level permissions*.

For actions that don't support resource-level permissions, such as listing operations, 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 <u>Resource types defined by</u> <u>Amazon Verified Permissions</u> 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 (or Condition *block*) lets you specify conditions in which a statement is in effect. The Condition element is optional. You can create conditional expressions that use <u>condition operators</u>, such as equals or less than, to match the condition in the policy with values in the request.

How Amazon Verified Permissions works with IAM

If you specify multiple Condition elements in a statement, or multiple keys in a single Condition element, AWS evaluates them using a logical AND operation. If you specify multiple values for a single condition key, AWS evaluates the condition using a logical OR operation. All of the conditions must be met before the statement's permissions are granted.

You can also use placeholder variables when you specify conditions. For example, you can grant an IAM user permission to access a resource only if it is tagged with their IAM user name. For more information, see IAM policy elements: variables and tags in the *IAM User Guide*.

AWS supports global condition keys and service-specific condition keys. 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. In AWS, these attributes are called *tags*. You can attach tags to IAM entities (users or roles) and to many AWS resources. Tagging entities and resources is the first step of ABAC. Then you design ABAC policies to allow operations when the principal's tag matches the tag on the resource that they are trying to access.

ABAC is helpful in environments that are growing rapidly and helps with situations where policy management becomes cumbersome.

To control access based on tags, you provide tag information in the <u>condition element</u> 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 <u>Define permissions with ABAC authorization</u> in the *IAM User Guide*. To view a tutorial with steps for setting up ABAC, see <u>Use attribute-based access control</u> (ABAC) in the *IAM User Guide*.

Using temporary credentials with Verified Permissions

Supports temporary credentials Yes

Some AWS services don't work when you sign in using temporary credentials. For additional information, including which AWS services work with temporary credentials, see <u>AWS services that</u> work with IAM in the *IAM User Guide*.

You are using temporary credentials if you sign in to the AWS Management Console using any method except a user name and password. For example, when you access AWS using your company's single sign-on (SSO) link, that process automatically creates temporary credentials. You also automatically create temporary credentials when you sign in to the console as a user and then switch roles. For more information about switching roles, see <u>Switch from a user to an IAM role</u> (console) in the *IAM User Guide*.

You can manually create temporary credentials using the AWS CLI or AWS API. You can then use those temporary credentials to access AWS. AWS recommends that you dynamically generate temporary credentials instead of using long-term access keys. For more information, see <u>Temporary security credentials in IAM</u>.

Cross-service principal permissions for Verified Permissions

Supports principal permissions Yes

When you use an IAM user or role to perform actions in AWS, you are considered a principal. When you use some services, you might perform an action that then initiates another action in a different service. FAS uses the permissions of the principal calling an AWS service, combined with the requesting AWS service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other AWS services or resources to complete. In this case, you must have permissions to perform both actions. 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 <u>IAM role</u> 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 <u>Create a role to delegate permissions to an AWS service</u> 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 <u>AWS services that work with IAM</u>. 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 <u>IAM JSON policy elements</u> reference in the IAM User Guide.

Action	Description	
CreateIdentitySource	Action to create a new identity source.	
<u>CreatePolicy</u>	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.	

Action	Description
IsAuthorized	Action to get an <u>authorization response</u> based on the parameters described in the <u>authoriza</u> <u>tion request</u> .
<u>IsAuthorizedWithToken</u>	Action to get an <u>authorization response</u> based on the parameters described in the <u>authoriza</u> <u>tion request</u> 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.
<u>ListPolicyTemplates</u>	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 <u>Creating IAM policies</u> 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 <u>Actions, resources, and condition keys for</u> <u>Amazon Verified Permissions</u> in the *Service Authorization Reference*.

Topics

- Policy best practices
- Using the Verified Permissions console
- <u>Allow users to view their own permissions</u>

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 <u>AWS managed policies</u> or <u>AWS</u> 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 <u>Policies and permissions in IAM</u> 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 AWS 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 <u>Validate policies with IAM Access Analyzer</u> 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 <u>Secure API</u> access with MFA in the IAM User Guide.

For more information about best practices in IAM, see <u>Security best practices in IAM</u> 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 <u>Adding permissions to a user</u> 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 <u>create IAM customer managed policies</u> 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 <u>AWS managed policies</u> 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 <u>AWS managed</u> <u>policies for job functions</u> 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.

(i) 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/*"
      1
    }
  1
}
```

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
<u>AmazonVerifiedPermissionsFu</u> <u>llAccess</u> – New policy	Verified Permissions added a new policy to allow full access to Verified Permissions.	October 11, 2024
AmazonVerifiedPermissionsRe adOnlyAccess – 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 IsAuthorizedWithTo ken .	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
- <u>I am not authorized to perform iam:Pas</u>sRole

• 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 *myexample-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 <u>How Amazon Verified</u>
 <u>Permissions works with IAM</u>.
- To learn how to provide access to your resources across AWS accounts that you own, see <u>Providing access to an IAM user in another AWS account that you own</u> in the *IAM User Guide*.
- To learn how to provide access to your resources to third-party AWS accounts, see <u>Providing</u> access to AWS accounts owned by third parties in the *IAM User Guide*.
- To learn how to provide access through identity federation, see <u>Providing access to externally</u> 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 <u>Cross account resource access in IAM</u> in the *IAM User Guide*.

Compliance validation for Amazon Verified Permissions

To learn whether an AWS service is within the scope of specific compliance programs, see <u>AWS</u> <u>services in Scope by Compliance Program</u> 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. AWS provides the following resources to help with compliance:

- <u>Security Compliance & Governance</u> These solution implementation guides discuss architectural considerations and provide steps for deploying security and compliance features.
- <u>HIPAA Eligible Services Reference</u> Lists HIPAA eligible services. Not all AWS services are HIPAA eligible.
- <u>AWS Compliance Resources</u> This collection of workbooks and guides might apply to your industry and location.
- <u>AWS Customer Compliance Guides</u> Understand the shared responsibility model through the lens of compliance. The guides summarize the best practices for securing AWS services and map the guidance to security controls across multiple frameworks (including National Institute of Standards and Technology (NIST), Payment Card Industry Security Standards Council (PCI), and International Organization for Standardization (ISO)).
- <u>Evaluating Resources with Rules</u> in the *AWS Config Developer Guide* The AWS Config service assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.
- <u>AWS Security Hub</u> This AWS service provides a comprehensive view of your security state within AWS. Security Hub uses security controls to evaluate your AWS resources and to check your compliance against security industry standards and best practices. For a list of supported services and controls, see <u>Security Hub controls reference</u>.
- <u>Amazon GuardDuty</u> This AWS service detects potential threats to your AWS accounts, workloads, containers, and data by monitoring your environment for suspicious and malicious activities. GuardDuty can help you address various compliance requirements, like PCI DSS, by meeting intrusion detection requirements mandated by certain compliance frameworks.
- <u>AWS Audit Manager</u> This AWS service helps you continuously audit your AWS usage to simplify how you manage risk and compliance with regulations and industry standards.

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 <u>AWS Global Infrastructure</u>.

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 <u>AWS CloudTrail User Guide</u>.

For more information about monitoring Verified Permissions with CloudTrail, see <u>Logging Amazon</u> 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 <u>AWS CloudTrail User Guide</u>.

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
- <u>CloudTrail supported services and integrations</u>
- Configuring Amazon SNS notifications for CloudTrail
- <u>Receiving CloudTrail log files from multiple regions</u> and <u>Receiving CloudTrail log files from</u> <u>multiple accounts</u>

All Verified Permissions actions are logged by CloudTrail and are documented in the <u>Amazon</u> <u>Verified Permissions API Reference Guide</u>. 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 <u>CloudTrail userIdentity element</u>.

Data events like <u>IsAuthorized</u> and <u>IsAuthorizedWithToken</u> 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 <u>Data events</u> 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

- IsAuthorized
- BatchIsAuthorized
- CreatePolicyStore
- ListPolicyStores
- DeletePolicyStore
- PutSchema
- GetSchema
- <u>CreatePolicyTemplate</u>
- <u>DeletePolicyTemplate</u>
- CreatePolicy
- GetPolicy
- <u>CreateIdentitySource</u>
- GetIdentitySource
- ListIdentitySources
- DeleteIdentitySource

1 Note

Some fields have been redacted from the examples for data privacy.

IsAuthorized

```
{
    "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"
```

BatchIsAuthorized

```
{
    "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"
}
```

CreatePolicyStore

```
{
 "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-TOKEN111111",
    "validationSettings": {
```

```
User Guide
```

```
"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

```
{
 "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

```
{
  "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

```
{
  "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

```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "EXAMPLE_PRINCIPAL_ID",
    "arn": "arn:aws:iam::222222222222:role/ExampleRole",
    "accountId": "22222222222",
    "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": "alf4d4cd-6156-480a-a9b8-e85a71dcc7c2",
  "eventID": "0b3b8e3d-155c-46f3-a303-7e9e8b5f606b",
  "readOnly": true,
  "resources": [
    {
      "accountId": "22222222222",
      "type": "AWS::VerifiedPermissions::PolicyStore",
      "ARN": "arn:aws:verifiedpermissions::222222222222:policy-store/
PSEXAMPLEabcdefg111111"
    }
  ],
  "eventType": "AwsApiCall",
  "managementEvent": true,
  "recipientAccountId": "22222222222",
  "eventCategory": "Management"
}
```

CreatePolicyTemplate

```
{
    "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

```
{
    "eventVersion": "1.08",
    "userIdentity": {
```

```
"type": "AssumedRole",
    "principalId": "EXAMPLE_PRINCIPAL_ID",
    "arn": "arn:aws:iam::222222222222:role/ExampleRole",
    "accountId": "22222222222",
    "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": "22222222222",
      "type": "AWS::VerifiedPermissions::PolicyStore",
      "ARN": "arn:aws:verifiedpermissions::222222222222;policy-store/
PSEXAMPLEabcdefg111111"
    }
  ],
  "eventType": "AwsApiCall",
  "managementEvent": true,
  "recipientAccountId": "22222222222",
  "eventCategory": "Management"
}
```

CreatePolicy

```
{
    "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-TOKEN111111",
    "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

```
{
  "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

{

```
"eventVersion": "1.08",
```

```
"userIdentity": {
    "type": "AssumedRole",
    "principalId": "EXAMPLE_PRINCIPAL_ID",
    "arn": "arn:aws:iam::3333333333333:role/ExampleRole",
    "accountId": "3333333333333",
    "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-TOKEN111111",
    "configuration": {
      "cognitoUserPoolConfiguration": {
        "userPoolArn": "arn:aws:cognito-idp:000011112222:us-east-1:userpool/us-
east-1_aaaaaaaaaaa"
      }
    },
    "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": "3333333333333",
      "type": "AWS::VerifiedPermissions::PolicyStore",
      "arn": "arn:aws:verifiedpermissions::3333333333333:policy-store/
PSEXAMPLEabcdefg111111"
    }
  ],
  "eventType": "AwsApiCall",
  "managementEvent": true,
  "recipientAccountId": "3333333333333",
```

}

{

```
"eventCategory": "Management"
```

GetIdentitySource

```
"eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "EXAMPLE_PRINCIPAL_ID",
    "arn": "arn:aws:iam::3333333333333:role/ExampleRole",
    "accountId": "3333333333333",
    "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": "3333333333333",
      "type": "AWS::VerifiedPermissions::PolicyStore",
      "arn": "arn:aws:verifiedpermissions::3333333333333:policy-store/
PSEXAMPLEabcdefg111111"
    }
  ],
  "eventType": "AwsApiCall",
  "managementEvent": true,
  "recipientAccountId": "333333333333",
  "eventCategory": "Management"
}
```

ListIdentitySources

```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "EXAMPLE_PRINCIPAL_ID",
    "arn": "arn:aws:iam::3333333333333:role/ExampleRole",
    "accountId": "3333333333333",
    "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": "3333333333333",
      "type": "AWS::VerifiedPermissions::PolicyStore",
      "arn": "arn:aws:verifiedpermissions::3333333333333:policy-store/
PSEXAMPLEabcdefg111111"
    }
  ],
  "eventType": "AwsApiCall",
  "managementEvent": true,
  "recipientAccountId": "333333333333",
  "eventCategory": "Management"
}
```

DeleteIdentitySource

```
{
    "eventVersion": "1.08",
    "userIdentity": {
```

```
"type": "AssumedRole",
    "principalId": "EXAMPLE_PRINCIPAL_ID",
    "arn": "arn:aws:iam::333333333333:role/ExampleRole",
    "accountId": "3333333333333",
    "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": "3333333333333",
      "type": "AWS::VerifiedPermissions::PolicyStore",
      "arn": "arn:aws:verifiedpermissions::3333333333333;policy-store/
PSEXAMPLEabcdefg111111"
    }
  ],
  "eventType": "AwsApiCall",
  "managementEvent": true,
  "recipientAccountId": "3333333333333",
  "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 AWS CloudFormation provisions and configures those resources for you.

When you use AWS 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.

<u> Important</u>

Amazon Cognito Identity is not available in all of the same AWS Regions as Amazon Verified Permissions. If you receive an error from AWS 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 AWS CloudFormation templates

To provision and configure resources for Verified Permissions and related services, you must understand <u>AWS CloudFormation templates</u>. Templates are formatted text files in JSON or YAML. These templates describe the resources that you want to provision in your AWS CloudFormation stacks. If you're unfamiliar with JSON or YAML, you can use AWS CloudFormation Designer to help you get started with AWS CloudFormation templates. For more information, see <u>What is AWS</u> <u>CloudFormation Designer?</u> in the *AWS CloudFormation User Guide*.

Verified Permissions supports creating identity sources, policies, policy stores, and policy templates in AWS CloudFormation. For more information, including examples of JSON and YAML templates for Verified Permissions resources, see the <u>Amazon Verified Permissions resource type reference</u> 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 AWS CloudFormation. Constructs, or reusable cloud components, can be used to create AWS CloudFormation templates. These templates can then be used to deploy your cloud infrastructure.

To learn more and download AWS CDK, see <u>AWS Cloud Development Kit</u>.

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

Amazon Verified Permissions L2 CDK Construct

Learn more about AWS CloudFormation

To learn more about AWS CloudFormation, see the following resources:

- AWS CloudFormation
- AWS CloudFormation User Guide
- AWS 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 AWS 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 <u>Access AWS services through AWS PrivateLink</u> in the AWS PrivateLink Guide.

Considerations for Verified Permissions

Before you set up an interface endpoint for Verified Permissions, review <u>Considerations</u> 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 <u>Create an interface</u> 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 <u>Control access to services using endpoint policies</u> 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 <u>Service Quotas console</u>. In the navigation pane, choose **AWS services** and select **Verified Permissions**.

To request a quota increase, see <u>Requesting a Quota Increase</u> in the *Service Quotas User Guide*. If the quota is not yet available in Service Quotas, use the <u>limit increase form</u>.

Your AWS account has the following quotas related to Verified Permissions.

Topics

- Quotas for resources
- Quotas for hierarchies
- Quotas for operations per second

Quotas for resources

Name	Default	Adjus e	Description
Policy stores per Region per account	Each supported Region: 30,000	<u>Yes</u>	The maximum number of policy stores.
Policy templates per policy store	Each supported Region: 40	<u>Yes</u>	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.

Amazon Verified Permissions

Name	Default	Adjus e	Description
Authorization request size ¹	1 MB	No	The maximum size of an authorization request.
Policy size	10,000 bytes	No	The maximum size of an individual policy.
Schema size	100,000 bytes	No	The maximum size of the schema of a policy store.
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 <u>IsAuthorized</u> and <u>IsAuthorizedWithToken</u>.

² The default limit for the total size of all the policies scoped for a single resource is 200,000 bytes. Similarly, the total size of all the policies, where the scope leaves the resource undefined, thereby applying to all resources, is limited by default to 200,000 bytes. Note that 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. This limit can be raised, provided that your policy design meets certain constraints. If you need to explore this option, <u>contact Support</u>.

Template-linked policy size example

You can determine how template-linked policies contribute to the *Policy size per resource* quota by taking the sum of the length of the principal and resource. If the principal or resource isn't specified, the length of that piece is 0. If a resource isn't specified, its size counts towards the "unspecified" resource quota. The size of the template body itself has no impact on the policy size.

Let's look at the following template:

```
@id("template1")
```

```
permit (
    principal in ?principal,
    action in [Action::"view", Action::"comment"],
    resource in ?resource
)
unless {
    resource.tag =="private"
};
```

Let's 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
}
```

Now, let's calculate the size of those policies by counting the characters in the principal and resource for each one. Each character counts as 1 byte.

The size of policy1 would be the length of the principal User::"alice" (13) plus the length of the resource Photo::"car.jpg" (16). Adding them up we have 13 + 16 = 29 bytes.

The size of policy2 would be the length of the principal User::"bob" (11) plus the length of the resource Photo::"boat.jpg" (17). Adding them up we have 11 + 17 = 28 bytes.

The size of policy3 would be the length of the principal User::"jane" (12) plus the length of the resource Photo::"car.jpg" (16). Adding them up we have 12 + 16 = 28 bytes.

The size of policy4 would be the length of the principal User::"jane" (12) plus the length of the resource (0). Adding them up we have 12 + 0 = 12 bytes.

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

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

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

Quotas for hierarchies


í) Note

The following quotas are aggregated, meaning they are added together. The maximum number of transitive parents for the group is what's listed. For example, if the limit of *Transitive parents per principal* is 100 that means there could be 100 parents of *principals* and 0 parents for both *actions* and *resources*, or any combination of parents that add up to 100 **total** parents.

Name	Default	Adjus e	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.

Name	Default	Adjus e	Description
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 <u>Service Quotas</u>. To prevent applications from exceeding the quota for an operation, you must optimize them for retries and exponential backoff. For more information, see <u>Retry with backoff pattern</u> and <u>Managing and monitoring API throttling in your workloads</u>.

Name	Default	Adjus e	Description
BatchGetPolicy requests per second per Region per account	Each supported Region: 10	<u>Yes</u>	The maximum number of BatchGetPolicy requests per second.
BatchIsAuthorized requests per second per Region per account	Each supported Region: 30	<u>Yes</u>	The maximum number of BatchIsAuthorized requests per second.
BatchIsAuthorizedWithToken requests per second per Region per account	Each supported Region: 30	<u>Yes</u>	The maximum number of BatchIsAuthorizedW ithToken requests per second.
CreateldentitySource requests per second per Region per account	Each supported Region: 1	<u>Yes</u>	The maximum number of CreateIdentitySource requests per second.
CreatePolicy requests per second per Region per account	Each supported Region: 10	<u>Yes</u>	The maximum number of CreatePolicy requests per second.
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 account	Each supported Region: 10	<u>Yes</u>	The maximum number of CreatePolicyTemplate requests per second.
DeleteIdentitySource requests per second per Region per account	Each supported Region: 1	<u>Yes</u>	The maximum number of DeleteIdentitySource requests per second.

Amazon Verified Permissions

Name	Default	Adjus e	Description
DeletePolicy requests per second per Region per account	Each supported Region: 10	<u>Yes</u>	The maximum number of DeletePolicy requests per second.
DeletePolicyStore requests per second per Region per account	Each supported Region: 1	No	The maximum number of DeletePolicyStore requests per second.
DeletePolicyTemplate requests per second per Region per account	Each supported Region: 10	<u>Yes</u>	The maximum number of DeletePolicyTemplate requests per second.
GetIdentitySource requests per second per Region per account	Each supported Region: 10	<u>Yes</u>	The maximum number of GetIdentitySource requests per second.
GetPolicy requests per second per Region per account	Each supported Region: 10	<u>Yes</u>	The maximum number of GetPolicy requests per second.
GetPolicyStore requests per second per Region per account	Each supported Region: 10	<u>Yes</u>	The maximum number of GetPolicyStore requests per second.
GetPolicyTemplate requests per second per Region per account	Each supported Region: 10	<u>Yes</u>	The maximum number of GetPolicyTemplate requests per second.
GetSchema requests per second per Region per account	Each supported Region: 10	<u>Yes</u>	The maximum number of GetSchema requests per second.
IsAuthorized requests per second per Region per account	Each supported Region: 200	<u>Yes</u>	The maximum number of IsAuthorized requests per second.

Amazon Verified Permissions

Name	Default	Adjus e	Description
lsAuthorizedWithToken requests per second per Region per account	Each supported Region: 200	<u>Yes</u>	The maximum number of IsAuthorizedWithToken requests per second.
ListIdentitySources requests per second per Region per account	Each supported Region: 10	<u>Yes</u>	The maximum number of ListIdentitySources requests per second.
ListPolicies requests per second per Region per account	Each supported Region: 10	<u>Yes</u>	The maximum number of ListPolicies requests per second.
ListPolicyStores requests per second per Region per account	Each supported Region: 10	<u>Yes</u>	The maximum number of ListPolicyStores requests per second.
ListPolicyTemplates requests per second per Region per account	Each supported Region: 10	<u>Yes</u>	The maximum number of ListPolicyTemplates requests per second.
PutSchema requests per second per Region per account	Each supported Region: 10	<u>Yes</u>	The maximum number of PutSchema requests per second.
UpdateIdentitySource requests per second per Region per account	Each supported Region: 1	<u>Yes</u>	The maximum number of UpdateIdentitySource requests per second.
UpdatePolicy requests per second per Region per account	Each supported Region: 10	<u>Yes</u>	The maximum number of UpdatePolicy requests per second.
UpdatePolicyStore requests per second per Region per account	Each supported Region: 10	No	The maximum number of UpdatePolicyStore requests per second.

Name	Default	Adjus e	Description
UpdatePolicyTemplate requests per second per Region per account	Each supported Region: 10	<u>Yes</u>	The maximum number of UpdatePolicyTemplate requests per second.

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
- <u>Authorization request</u>
- <u>Authorization response</u>
- Considered policies
- Context data
- Determining policies
- Entity data
- Permissions, authorization, and principals
- Policy enforcement
- Policy store
- Satisfied policies
- Differences between Amazon Verified Permissions and the Cedar policy language

Cedar policy language concepts

- Authorization
- Entity
- Groups and hierarchies
- <u>Namespaces</u>
- Policy
- Policy template
- <u>Schema</u>

Authorization model

The *authorization model* describes the scope of the <u>authorization requests</u> 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 <u>authorization request</u>. 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 <u>authorization request</u>.

Context data

Context data are attribute values that provide additional information to be evaluated.

Determining policies

Determining policies are the policies that determine the <u>authorization response</u>. For example, if there are two <u>satisfied policies</u>, 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 <u>policies</u>.

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 <u>authorization request</u>, 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.

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.

Namespace definition

Verified Permissions implementation of Cedar has the following differences from the native Cedar implementation:

- Verified Permissions supports only one <u>namespace in a schema</u> defined in a policy store.
- Verified Permissions doesn't allow you to create a <u>namespace</u> that's an empty string or includes the following values: aws, amazon, or cedar.

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

]

```
Ε
    {
        "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

Amazon Verified Permissions

Object	Verified Permissions limit (in bytes)	Cedar limit
Policy template size	10,000	None
Schema size	100,000	None
Entity type	200	None
Policy ID	64	None
Policy template ID	64	None
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 v4 FAQ

Amazon Verified Permissions is in the process of updating to Cedar v4. We're working to make this as seamless as possible for you. The following FAQs should answer your questions and help you prepare.

Topics

- What is the current state on the upgrade?
- Do I need to do anything right now?
- Does the upgrade of the console impact the authorization service?
- What are the breaking changes in Cedar v3 and Cedar v4?
- When will the upgrade to Cedar v4 be complete?

What is the current state on the upgrade?

As a first step we've upgraded the console to use Cedar v4.3, however the back end is still running on Cedar v2.5.0. This means that while you can now use the console to author policies using new features like the is operator, when you attempt to save them you will still get an error until we complete the upgrade.

Do I need to do anything right now?

No. You can start exploring Cedar v4 using the console, if you want, but don't need to do anything.

Does the upgrade of the console impact the authorization service?

No. Prior to upgrading we will run tests to check that your policy store works correctly with Cedar v4. There are some minor breaking changes between v2.5.0 and v4.3, but it's very unlikely that your policy store will be impacted. If it is, then your policy store will not be upgraded, and will continue to authorize using Cedar v2.5.0. Should this happen, we will reach out to explain any changes you need to make before being able to upgrade.

What are the breaking changes in Cedar v3 and Cedar v4?

Breaking changes are identified in the Cedar change log, marked with an (*).

🚯 Note

If your policy store is affected by breaking changes it won't be upgraded, and we'll work with you to update the policy store so it can be upgraded.

When will the upgrade to Cedar v4 be complete?

Our goal is for all accounts to be upgraded by December 31, 2025.

Document history for the Amazon Verified Permissions User Guide

The following table describes the documentation releases for Verified Permissions.

Change	Description	Date
New AWS managed policies	You can now use the AmazonVerifiedPerm issionsFullAccess and AmazonVerifiedPerm issionsReadOnlyAcc ess 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 BatchIsAu thorizedWithToken API request.	April 5, 2024
<u>Creating a policy store with</u> <u>API Gateway</u>	You can now create a policy store from an existing API and Amazon Cognito user pool.	April 1, 2024
<u>Context concepts and</u> <u>example</u>	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 AWS CloudFormation.	June 30, 2023
Initial release	Initial release of the Amazon Verified Permissions User Guide	June 13, 2023