



Resilient Architecture Guide

Amazon Pinpoint



Amazon Pinpoint: Resilient Architecture Guide

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

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

Table of Contents

.....	v
Welcome	1
What makes a workload "resilient"?	1
About Amazon Pinpoint	2
Regional availability	2
Additional resources	2
Resilient architecture overview	3
Is a resilient architecture necessary?	3
High availability concepts	3
When to fail over	4
Best practices	5
Enable the event stream	5
Use exponential backoff	5
Tailor your retry strategy to the channels that you use	6
Replicate infrastructure with code	6
Channel considerations	8
Email channel	8
Reputation	8
Opt-outs	9
Templates	9
SMS channel	10
Country-specific regulations	10
Fallback numbers	10
Throughput limits	11
Opt-outs	11
Templates	11
Other channels	12
Push notification channel	12
Custom channels	12
Voice channel	13
In-app notification channel	13
Synchronizing customer data	15
Endpoints	15
Create and update individual endpoints	15

Create or update endpoints in bulk	16
Events	18
Replicating event data	18
Reference architectures	20
Active-active	20
Active-active architecture considerations	20
Active-active architecture details	21
Benefits and drawbacks of an active-active architecture	23
Warm standby	23
Warm standby architecture considerations	24
Warm standby architecture details	24
Benefits and drawbacks of a warm standby architecture	26
Security	27
Data protection	27
Data encryption	30
Internetwork traffic privacy	30
Identity and access management	31
Audience	32
Authenticating with identities	32
Managing access using policies	33
How Amazon Pinpoint works with IAM	35
Amazon Pinpoint policy actions	41
Identity-based policy examples	79
IAM roles for common tasks	93
Troubleshooting	111
Security event logging and monitoring	113
Compliance validation	115
Infrastructure security	116
Security best practices	117
Document history	118

End of support notice: On October 30, 2026, AWS will end support for Amazon Pinpoint. After October 30, 2026, you will no longer be able to access the Amazon Pinpoint console or Amazon Pinpoint resources (endpoints, segments, campaigns, journeys, and analytics). For more information, see [Amazon Pinpoint end of support](#). **Note:** APIs related to SMS, voice, mobile push, OTP, and phone number validate are not impacted by this change and are supported by AWS End User Messaging.

Welcome to the Amazon Pinpoint Resilient Architecture Guide

Welcome to the *Amazon Pinpoint Resilient Architecture Guide*. This guide describes important factors to consider when designing a multi-Region architecture for Amazon Pinpoint. It also includes example reference architectures.

Many Amazon Pinpoint customers use Amazon Pinpoint to run mission-critical applications and services. For use cases that can only tolerate minimal downtime, it's important to consider deploying your workloads across several AWS Regions.

The architectures described in this guide might not fit every use case. Determining the right design for your needs requires an understanding of your application's business criticality, dependencies, workload volumes, and the nature of the work that it performs. Also keep in mind that this guide only considers multi-Region designs for Amazon Pinpoint. If your AWS use case includes other services, you should consider those workloads when you design a resilient architecture. If your organization has an AWS Enterprise Support plan, we recommend that you work closely with your Account Manager and Solutions Architects to develop an architecture that meets your specific requirements.

For information about using the features of Amazon Pinpoint through the AWS Management Console, see the [Amazon Pinpoint User Guide](#).

For information related to integrating Amazon Pinpoint with your web and mobile applications, see the [Amazon Pinpoint Developer Guide](#).

For reference content related to the Amazon Pinpoint API, see the [Amazon Pinpoint API Reference](#).

What makes a workload "resilient"?

A resilient workload is able to recover when stressed by load (more requests for service), attacks (either accidentally through a bug, or deliberately), and the failure of any component in the workload.

For more information about resiliency, see [Resiliency](#) on the AWS Well-Architected Framework website.

About Amazon Pinpoint

Amazon Pinpoint is an AWS service that you can use to engage with your customers across multiple messaging channels. You can use Amazon Pinpoint to send push notifications, in-app notifications, emails, text messages, voice messages, and messages over custom channels. It includes segmentation, campaign, and journey features that can help you send the right message to the right customer at the right time over the right channel.

Regional availability

Amazon Pinpoint is available in several AWS Regions in North America, Europe, Asia, and Oceania. In each Region, AWS maintains multiple Availability Zones. These Availability Zones are physically isolated from each other, but they are united by private, low-latency, high-throughput, and highly redundant network connections. These Availability Zones enable AWS to provide high levels of availability and redundancy, while also minimizing latency.

To learn more about AWS Regions, see [Specify which AWS Regions your account can use](#) in the *Amazon Web Services General Reference*.

For a list of all the Regions where Amazon Pinpoint is currently available and the endpoint for each Region, see [Amazon Pinpoint endpoints and quotas](#) and [AWS service endpoints](#) in the *Amazon Web Services General Reference*.

To learn more about the number of Availability Zones that are available in each Region, see [AWS global infrastructure](#).

Additional resources

AWS provides several additional resources for designing resilient architectures. For an overview of disaster recovery at AWS and related concepts, see the [AWS Disaster Recovery](#) Workshop and the [Reliability Pillar](#) section in the [Well-Architected Framework documentation](#).

Resilient architecture overview

This chapter contains introductory information related to the development of resilient, high availability architectures. It includes terms, concepts, and best practices.

Topics in this chapter:

- [Is a resilient, multi-Region architecture necessary for your use case?](#)
- [High availability concepts](#)
- [When to fail over to another AWS Region](#)

Is a resilient, multi-Region architecture necessary for your use case?

AWS was designed to help you achieve your system availability goals. Even if you only deploy services in a single AWS Region, those services are distributed across several Availability Zones in that Region. The result is high availability, geographical redundancy, and fault tolerance.

For critical use cases, consider using a resilient, multi-Region architecture. The primary benefit of using a multi-Region architecture is that it protects you against disruptions that impact an entire AWS Region. However, deploying this type of architecture requires you to make deeper investments in building your applications and regularly testing your failover capabilities. Weigh these benefits and drawbacks carefully against the criticality of your use case.

For more information about AWS Regions and Availability Zones, see [Regions and Availability Zones](#) on the AWS Global Infrastructure website.

High availability concepts

This guide uses several terms to describe common concepts in high availability architecture:

Recovery Time Objective (RTO)

The maximum acceptable delay between service interruption and service restoration. RTO determines what is considered an acceptable amount of time for the service to be unavailable.

Recovery Point Objective (RPO)

The maximum acceptable time since the last data recovery point. RPO determines what is considered an acceptable loss of data between the last recovery point and the service outage.

Warm standby

A high availability architecture in which a fully functional environment is always running in a secondary AWS Region. Business-critical systems are fully duplicated and are always on. If the primary Region becomes unavailable, you can use services such as Amazon Route 53 or AWS Global Accelerator to route all user traffic to the standby Region. The RPO for this architecture is typically measured in seconds, and the RTO is typically measured in minutes.

Active-active

A high availability architecture in which a workload is deployed in and actively serves traffic from multiple AWS Regions. An active-active design requires you to synchronize users and data between the Regions that you use. If one Region becomes unavailable, you can use services such as Amazon Route 53 or AWS Global Accelerator to route all user traffic to the other Region. The RPO and RTO for this type of architecture are measured in seconds.

There are other high availability architecture strategies that aren't described here, such as *pilot light* and *backup and restore*. However, these strategies aren't preferable for architectures that use Amazon Pinpoint. For that reason, this guide focuses on [warm standby](#) and [active-active](#) architectures.

When to fail over to another AWS Region

Several factors could cause your architecture to fail over to a different AWS Region. For example, a Regional outage could prevent you from accessing the Amazon Pinpoint console, or from accessing its API operations. You could also configure your architecture to fail over when your messages are being sent but aren't receiving event notifications (or the number of event notifications is unexpectedly low).

In certain situations, failing over won't provide any benefit. For example, if you send SMS messages, and a specific mobile carrier is having an outage, then delivery issues will persist, regardless of which AWS Region you use. The same is true for email: if an email provider has a temporary issue that prevents the delivery of email to its domain, that issue will persist across Regions.

Best practices for creating resilient Amazon Pinpoint architectures

This chapter contains best practices for creating resilient architectures in Amazon Pinpoint.

Topics in this chapter:

- [Enable the Amazon Pinpoint event stream](#)
- [Use exponential backoff](#)
- [Tailor your retry strategy to the channels that you use](#)
- [Replicate infrastructure with code](#)

Enable the Amazon Pinpoint event stream

When you build a resilient, multi-Region architecture, you should always enable the Amazon Pinpoint event stream. The event stream automatically delivers important information about the campaigns that you send. It also provides data about the email and SMS messages that you send, including information about whether those messages were delivered. This information is useful for tracking message delivery rates, and for troubleshooting delivery issues. For more information about enabling the Amazon Pinpoint event stream, see [Streaming Amazon Pinpoint events to Kinesis](#) in the *Amazon Pinpoint Developer Guide*.

As you plan your resilient Amazon Pinpoint architecture, make sure that you account for the replication of event data in your resilient Amazon Pinpoint architecture. For more information about duplicating event data, see [Synchronizing event data](#).

Use exponential backoff

If your calls to an AWS API result in failure, we recommend that you wait an increasing amount of time before issuing the command again. For more information, see [Retry behavior](#) in the *AWS General Reference*.

Tailor your retry strategy to the channels that you use

Depending on the channels that you use with Amazon Pinpoint, you might need to adopt a different message retry strategy. Some channels, such as push notifications, provide results in real time. That is, either your message is accepted and sent to the push notification service, or it isn't.

The nature of other channels, such as SMS and email, is that delivery of your message might be delayed because of factors outside of the control of AWS. For example, when you send an SMS message, it might not be delivered right away because the recipient's device isn't on a mobile network, their device is turned off, or there's a temporary issue with the mobile network. In these cases, the recipient's mobile provider might attempt to redeliver the message for several hours, which increases the amount of time that passes before you receive a success or failure notification. Even if a message is delivered right away, it might still take several minutes for the mobile carrier to send an acknowledgment back to the Amazon Pinpoint event stream.

Make sure that you account for the asynchronous behavior of the email and SMS channels when designing your retry strategy. Don't expect delivery confirmations to arrive immediately. Configure your monitoring systems with this asynchronous behavior in mind.

Replicate infrastructure with code

We recommend that you replicate your infrastructure in other Regions in an automated way. By automating the replication process, you can minimize human error. AWS services, such as CloudFormation, make it possible to recreate your architecture in other AWS Regions. Having a repeatable way to deploy and update AWS services helps your configurations remain consistent across Regions.

Amazon Pinpoint supports CloudFormation. However, there are some Amazon Pinpoint capabilities that CloudFormation doesn't currently include, including the following:

- **Journeys** – Amazon Pinpoint Journeys can't currently be created by using CloudFormation templates. However, the Amazon Pinpoint API includes extensive support for [creating journeys](#), [updating journeys](#), [viewing the details of a journey](#), and [changing the state of a journey](#). You can use AWS Lambda functions to call the necessary Amazon Pinpoint API operations, and you can include those functions in your CloudFormation templates.
- **SMS origination identities** – Origination identities are the identities that are used for sending SMS messages. Examples of origination identities include short codes, long codes, toll-free numbers, and 10DLC numbers. In some countries, obtaining these origination identities

requires a registration process. For example, to obtain a Sender ID for sending SMS messages to recipients in India, you must register your company and use case with regulatory authorities. Along similar lines, if you want to use a short code to send SMS messages to recipients in the United States, you must register your use case with the mobile carriers.

For more information about requesting various types of origination identities, see [Requesting SMS support](#) in the *Amazon Pinpoint User Guide*. For more information about the types of origination identities that are available in each country, see [Supported countries and regions](#) in the *Amazon Pinpoint User Guide*.

Channel resiliency considerations in Amazon Pinpoint

There are several factors that you should consider when creating a resilient Amazon Pinpoint architecture. This section provides information about concerns that are specific to certain communication channels, such as email and SMS.

Topics in this chapter:

- [Email channel architecture considerations](#)
- [SMS channel architecture considerations](#)
- [Architecture considerations for other channels](#)

Email channel architecture considerations

When designing a multi-Region architecture for Amazon Pinpoint, make sure that you consider two important factors in your architecture: reputation and opt-outs.

Topics in this section:

- [Reputation](#)
- [Opt-outs](#)
- [Templates](#)

Reputation

When deciding whether to mark an email as spam or deliver it to the inbox, email providers calculate a reputation score. This score considers multiple factors to determine if the email you send is trustworthy. There are two aspects of reputation to consider:

- **IP reputation** – Each IP address that you use to send email has its own reputation score. Many customers, especially those who send high volumes of email, use dedicated IP addresses. These dedicated IP addresses are unique to each AWS Region, and they can't be used across Regions. If you use dedicated IPs, you must consider if you want to use dedicated IPs in all of the Regions in which you use Amazon Pinpoint. If you maintain dedicated IPs in multiple Regions, it's vital to send regular, predictable volumes of email through those IPs to keep them warm. Email providers don't like to see massive increases in email volume from IP addresses that were previously inactive.

For more information about requesting dedicated email IP addresses, see [Requesting and relinquishing dedicated IP addresses](#) in the *Amazon Pinpoint User Guide*.

- **Domain and subdomain reputation** – Email providers also calculate Reputation scores based on your sending domain, regardless of the IP addresses that you use to send your messages. If you use a multi-Region architecture with Amazon Pinpoint, you must configure DomainKeys Identified Mail (DKIM) and Sender Policy Framework (SPF) in each Region. This means that you must create multiple records in the DNS configuration for your domain.

For more information about verifying email domains, see [Verifying a domain](#) in the *Amazon Pinpoint User Guide*.

Opt-outs

When your recipients unsubscribe from email communications, you must manage that preference across all AWS Regions. Email providers will block messages from domains that continue to send messages after users opt out.

The Amazon Pinpoint [event stream](#) provides data about unsubscribe events. For more information about email events, see [Email events](#) in the *Amazon Pinpoint Developer Guide*.

There are various techniques for synchronizing email events across AWS Regions. The better technique depends on which destination that you send email events to. For example, if you send events to Amazon S3, you can use the Cross-Region Replication (CRR) feature to automatically duplicate data in another Region. For more information about synchronization, see [Synchronizing event data](#).

Templates

If you use email templates, you should have copies of those templates in each Region. You can duplicate email templates by using the [ListTemplates](#) API operation to find the ID of the template that you want to duplicate. When you have the template ID, you can use the [GetEmailTemplate](#) operation to obtain the details of the template. Finally, you can use the [CreateEmailTemplate](#) operation to create the template in another Region.

You can also use the `AWS::Pinpoint::EmailTemplate` entity in AWS CloudFormation to deploy email templates to other AWS Regions automatically. For more information, see [AWS::Pinpoint::EmailTemplate](#) in the *AWS CloudFormation User Guide*.

SMS channel architecture considerations

You can use Amazon Pinpoint to send SMS messages to recipients in almost every country around the world. When you design a multi-Region architecture for sending SMS messages, you must consider several factors: country-specific regulations, fallback numbers, throughput limits, and opt-outs.

Country-specific regulations

Setting up a system for sending SMS messages globally can be difficult and time-consuming. Governments, regulators, and mobile carriers in each country have their own rules about application-to-person SMS messaging, including messages sent from services such as Amazon Pinpoint.

For example, when you send messages to recipients in the United States, you must use a dedicated phone number, such as a short code, toll-free number, or 10DLC number. Phone numbers are resources that are unique to each AWS Region, so they can't be shared across Regions. Additionally, each of these phone number types has its own registration process and a different set of capabilities. If you require the ability to send more than 100 messages per second, you must [request short codes for messaging](#). Short codes can take several weeks or even months to obtain, and the carriers impose several requirements on them, so it's important to plan ahead. You can also use a combination of phone number types. For example, if you have a warm standby architecture, your active Region might use a short code, while your standby Regions have [10DLC numbers](#) (which support lower throughput rates than short codes).

Another example of a country with specific requirements is India. If your business is based in India, you can register an alphanumeric Sender ID for sending your messages. This Sender ID can identify your brand in your messages, and it's less expensive to send messages to Indian recipients using a registered Sender ID than through a non-registered phone number. If you use an Indian Sender ID in multiple AWS Regions, you must set up your account to use that Sender ID in all of those Regions. For more information, see [India Sender ID registration process](#).

Fallback numbers

Different types of phone numbers can use different routes to reach your recipients. For example, in the US, a short code takes a different path downstream from AWS to reach end users than a 10DLC number takes. Having a variety of options, both within the same AWS Region and across Regions, can provide additional resiliency.

Throughput limits

In an [active-active](#) architecture, you can use Amazon Pinpoint to have the same SMS throughput in all Regions. However, in a [warm standby](#) architecture, it might be acceptable to have slightly lower throughput rates in your standby Regions. With SMS messaging, it typically costs more to send at higher throughput rates because you must obtain resources such as short codes. As you design a multi-Region SMS architecture, make sure to evaluate your throughput requirements.

Opt-outs

When your recipients unsubscribe from SMS communications, you must manage that preference across all AWS Regions. In some countries, there can be severe monetary penalties for not honoring the SMS opt-out preferences of your recipients. Additionally, mobile carriers can block messages from your phone number or Sender ID if they determine that you aren't honoring opt-outs.

The Amazon Pinpoint event stream emits data about opt-out events. For more information about SMS events, see [SMS events](#) in the *Amazon Pinpoint Developer Guide*.

There are various techniques for synchronizing SMS events across AWS Regions. The better technique depends on which destination that you send events to. For example, if you send events to Amazon S3, you can use the Cross-Region Replication (CRR) feature to duplicate data in another Region automatically. For more information about SMS events, see [Synchronizing event data](#).

Templates

If you use SMS templates, you should have copies of those templates in each Region. You can duplicate SMS templates by using the [ListTemplates](#) API operation to find the ID of the template that you want to duplicate. When you have the template ID, you can use the [GetSmsTemplate](#) operation to obtain the details of the template. Finally, you can use the [CreateSmsTemplate](#) operation to create the template in another Region.

You can also use the `AWS::Pinpoint::SmsTemplate` entity in AWS CloudFormation to deploy SMS templates to other AWS Regions automatically. For more information, see [AWS::Pinpoint::SmsTemplate](#) in the *AWS CloudFormation User Guide*.

Architecture considerations for other channels

Amazon Pinpoint supports several communication channels. This section includes factors that you should account for when designing a resilient architecture that includes the push notification, in-app notification, voice, or custom channels.

Other sections in this guide contain dedicated sections related to using the SMS and email channels in a multi-Region architecture. For more information about those channels, see [Email channel architecture considerations](#) and [SMS channel architecture considerations](#).

Push notification channel

If you use the push notification channel in multiple AWS Regions, you must separately configure the push notification services (such as Firebase Cloud Messaging or Apple Push Notification Service) in each Region. Typically, this involves providing an API key or certificate for each notification service.

For more information about setting up the push notification channel, see [Setting up Amazon Pinpoint mobile push channels](#) in the *Amazon Pinpoint User Guide*.

If you use templates for your push notifications, make sure that your templates exist in each Region. You can duplicate push notification templates by using the [ListTemplates](#) API operation to find the ID of the template that you want to duplicate. When you have the template ID, you can use the [GetPushTemplate](#) operation to obtain the details of the template, and then you can use the [CreatePushTemplate](#) operation to create the template in another Region.

Custom channels

Custom channels in Amazon Pinpoint can call AWS Lambda functions or webhooks. If you use custom channels in multiple AWS Regions, make sure that all of the resources that the custom channel relies on are present in each Region. For example, if your custom channel calls a Lambda function, that function must be present in each Region. If that Lambda function calls other AWS services, then you must reproduce those services in each Region.

For more information about custom channels, see [Creating custom channels in Amazon Pinpoint](#) in the *Amazon Pinpoint Developer Guide*.

You can use the `AWS::Lambda::Function` entity in AWS CloudFormation to deploy custom channel Lambda functions to other AWS Regions automatically. For more information, see [AWS::Lambda::Function](#) in the *AWS CloudFormation User Guide*.

Voice channel

To send voice messages to recipients in a specific country, you must have a dedicated phone number for that country. If you use the voice channel in multiple AWS Regions, you must obtain phone numbers in each Region. Phone numbers are Region-specific resources, which means that the same phone number can't be used in more than one AWS Region.

For more information about setting up the voice channel, see [Setting up the Amazon Pinpoint voice channel](#) in the *Amazon Pinpoint User Guide*.

By default, new Amazon Pinpoint accounts are placed in the voice sandbox. While an account is in the sandbox, Amazon Pinpoint places limits on the messages that you can send. For example, you can only send 20 messages per day, and you can only send messages to a limited number of countries. You can request to have your account removed from the sandbox, which removes these limitations. If you use the voice channel, make sure that you have your account removed from the sandbox in each Region.

For more information about setting the voice sandbox, see [Amazon Pinpoint voice sandbox](#) in the *Amazon Pinpoint User Guide*.

If you use templates for your voice messages, make sure that your templates exist in each Region. You can duplicate voice templates by using the [ListTemplates](#) API operation to find the ID of the template that you want to duplicate. When you have the template ID, you can use the [GetVoiceTemplate](#) operation to obtain the details of the template, and then use the [CreateVoiceTemplate](#) operation to create the template in another Region.

In-app notification channel

The in-app notification channel in Amazon Pinpoint relies on your campaigns or journeys creating in-app messages for a given endpoint, and it relies on your apps retrieving those messages from Amazon Pinpoint. If you use the in-app notification channel in multiple AWS Regions, make sure that your in-app message templates exist in each Region. You can duplicate in-app message templates by using the [ListTemplates](#) API operation to find the ID of the template that you want to duplicate. When you have the template ID, you can use the [GetInAppTemplate](#) operation to obtain the details of the template. Finally, you can use the [CreateInAppTemplate](#) operation to create the template in another Region.

You can also use the `AWS::Pinpoint::InAppTemplate` entity in AWS CloudFormation to deploy in-app message templates to other AWS Regions automatically. For more information about in-app message templates, see [AWS::Pinpoint::InAppTemplate](#) in the *AWS CloudFormation User Guide*.

For more information about the in-app channel, see [Sending and retrieving in-app messages in Amazon Pinpoint](#) in the *Amazon Pinpoint Developer Guide*.

Synchronizing customer data across AWS Regions

Regardless of the architecture design that you choose, make sure that your customer data is synchronized across the AWS Regions that you intend to use. In this case, "customer data" refers to the contact information for your customers (such as their email addresses, phone numbers, name, or company). It also refers to the preference data for your customers—that is, their opt-in and opt-out preferences. Finally, it refers to information about whether they're able to receive messages from you.

In a resilient architecture, it's important to keep all of this information synchronized across all of the AWS Regions in which you use Amazon Pinpoint. This chapter contains example architectures that you can use to keep this information synchronized.

Topics in this chapter:

- [Synchronizing endpoint information](#)
- [Synchronizing event data](#)

Synchronizing endpoint information

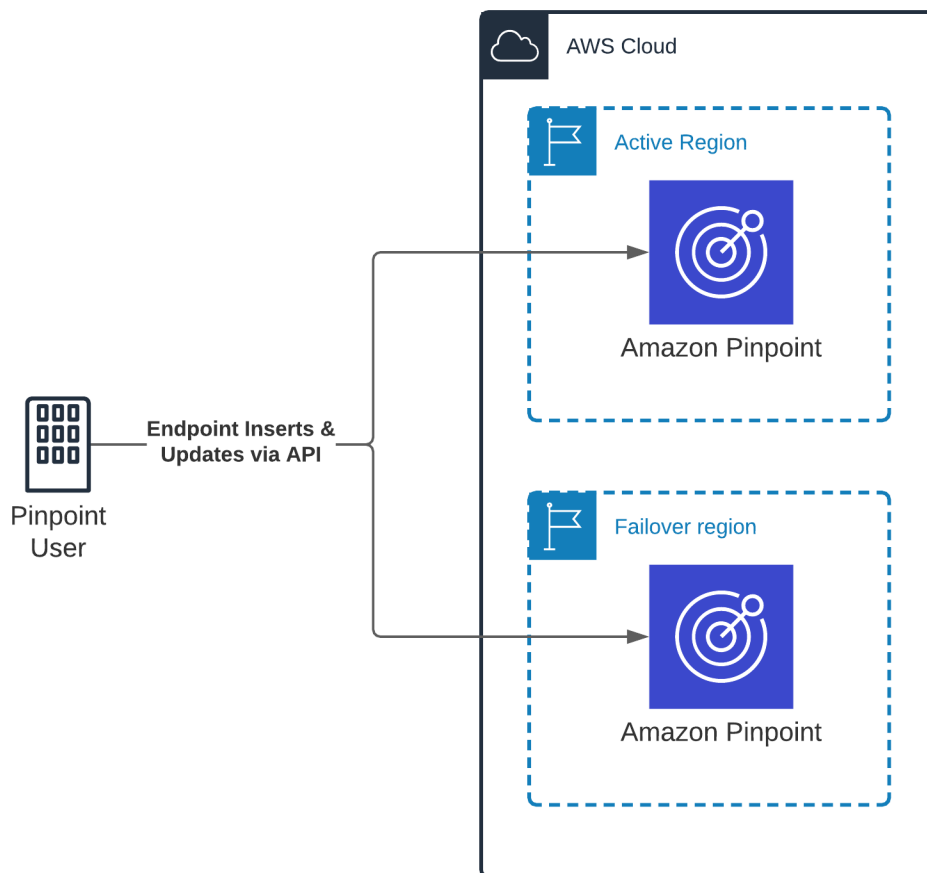
This section contains architecture examples related to synchronizing Amazon Pinpoint endpoint data across multiple AWS Regions.

Topics in this section:

- [Create and update individual endpoints](#)
- [Create or update endpoints in bulk](#)

Create and update individual endpoints

If your application or service uses the [UpdateEndpoint](#) API to create individual endpoints in your Amazon Pinpoint account, you can simultaneously call the API in each of your target Regions. Your call to each Region contains exactly the same data. This solution is suitable for situations in which you are adding new endpoints to Amazon Pinpoint as the endpoints contact information is being captured, instead of loading endpoints from an existing database or system in bulk. This approach works for both [active-active](#) and [warm standby](#) architectures, and is illustrated in the following image:

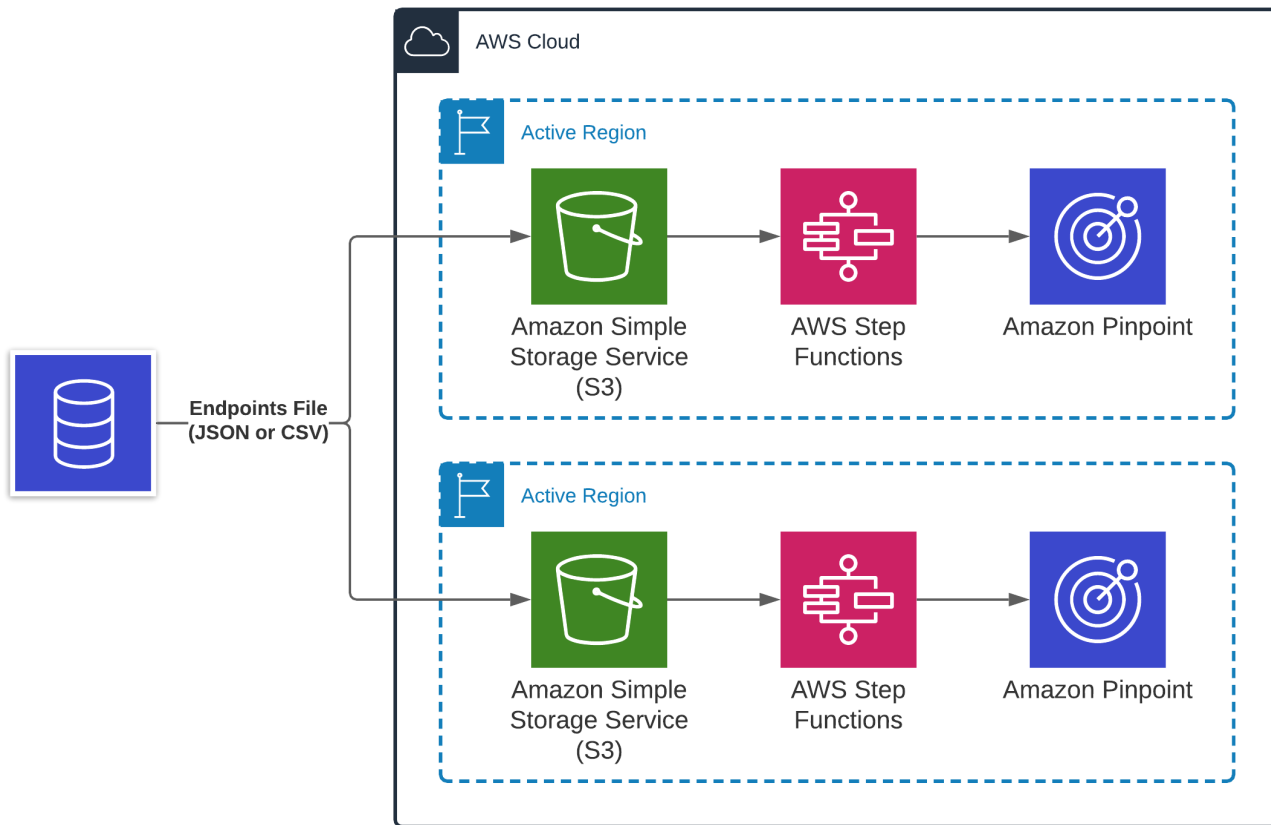


The benefit of this architecture is that the Recovery Point Objective (RPO) is near zero. A disadvantage is that you have to make twice as many API calls at ingestion time, which could increase latency when creating or updating endpoints.

Create or update endpoints in bulk


A common way that you can manage endpoints is to export customer data from a data lake or system of record. The exported data is then stored in an Amazon S3 bucket. An Import Job then picks up the data from the Amazon S3 bucket and imports it into Amazon Pinpoint. If your architecture uses Import Jobs, make sure that you import the data into all of your AWS Regions.

You can do these imports in parallel by writing the source files to Amazon S3 buckets in each of your target Regions. You can then use Lambda functions and Step Functions workflows to import that data in each Region automatically. This architecture works for both [active-active](#) and [warm standby](#) architectures. The parallel import architecture is illustrated in the following diagram:



The benefit of this architecture is that the RPO is near zero. A disadvantage is that you have to make twice as many API calls at ingestion time, which could increase latency when creating or updating endpoints.

If you use a warm standby architecture, you could alternatively use a recurring *Export Job*. The Export Job exports segment members into an Amazon S3 bucket in the warm standby Region. You can then use Lambda functions and Step Functions workflows to create Import Jobs in the warm standby Region. This architecture is illustrated in the following diagram:




A disadvantage of this architecture is that the RPO is increased. However, you can control the RPO by changing how often Export Jobs is performed.

Synchronizing event data

When you use Amazon Pinpoint to send email or SMS messages, it generates event records. For email, these event records can tell you whether a message was delivered, rejected, opened, clicked, and more. For SMS messages, event records can tell you whether a message was delivered. These records are useful for tracking the success of your messages, and for troubleshooting issues. This section contains information about synchronizing event data across AWS Regions.

Replicating event data

Amazon Pinpoint can send events to a Amazon Kinesis Data Firehose stream. The Firehose stream can then send that data to numerous destinations, including Amazon S3 buckets and Amazon Redshift clusters. Many of these destinations support the automatic replication of data across AWS Regions. For example, Amazon S3 includes a feature called Cross-Region Replication (CRR). The following diagram shows an example of an Amazon Pinpoint architecture that uses Amazon S3 CRR:



For more information about CRR, see [Replicating objects overview](#) in the *Amazon S3 User Guide*.

If you send your event data to an Amazon Redshift cluster instead of an Amazon S3 bucket, you can implement a similar architecture using cross-Region data sharing. For more information, see [Sharing data across AWS Regions](#) in the *Amazon Redshift Database Developer Guide*.

Amazon Pinpoint reference architectures

This chapter contains two sample architectures for resilient Amazon Pinpoint workloads: an active-active architecture, and a warm standby architecture. Each of these architectures has its own distinct benefits and drawbacks. The topics in this chapter provide information about each of these architecture types.

Topics in this chapter:

- [Amazon Pinpoint active-active reference architecture](#)
- [Amazon Pinpoint warm standby reference architecture](#)

Amazon Pinpoint active-active reference architecture

This section describes an active-active architecture for Amazon Pinpoint. In this type of architecture, two identical instances of Amazon Pinpoint are maintained in two separate AWS Regions. Messages are sent through both Regions simultaneously. If an increased rate of errors occurs in one Region, the traffic that would normally be sent through the affected Region is instead diverted to the other Region. When the error rate in the impacted Region returns to normal, the traffic is again divided between both Regions.

Topics in this section:

- [Active-active architecture considerations](#)
- [Active-active architecture details](#)
- [Benefits and drawbacks of an active-active architecture](#)

Active-active architecture considerations

Consider the following factors when implementing an active-active architecture with Amazon Pinpoint:

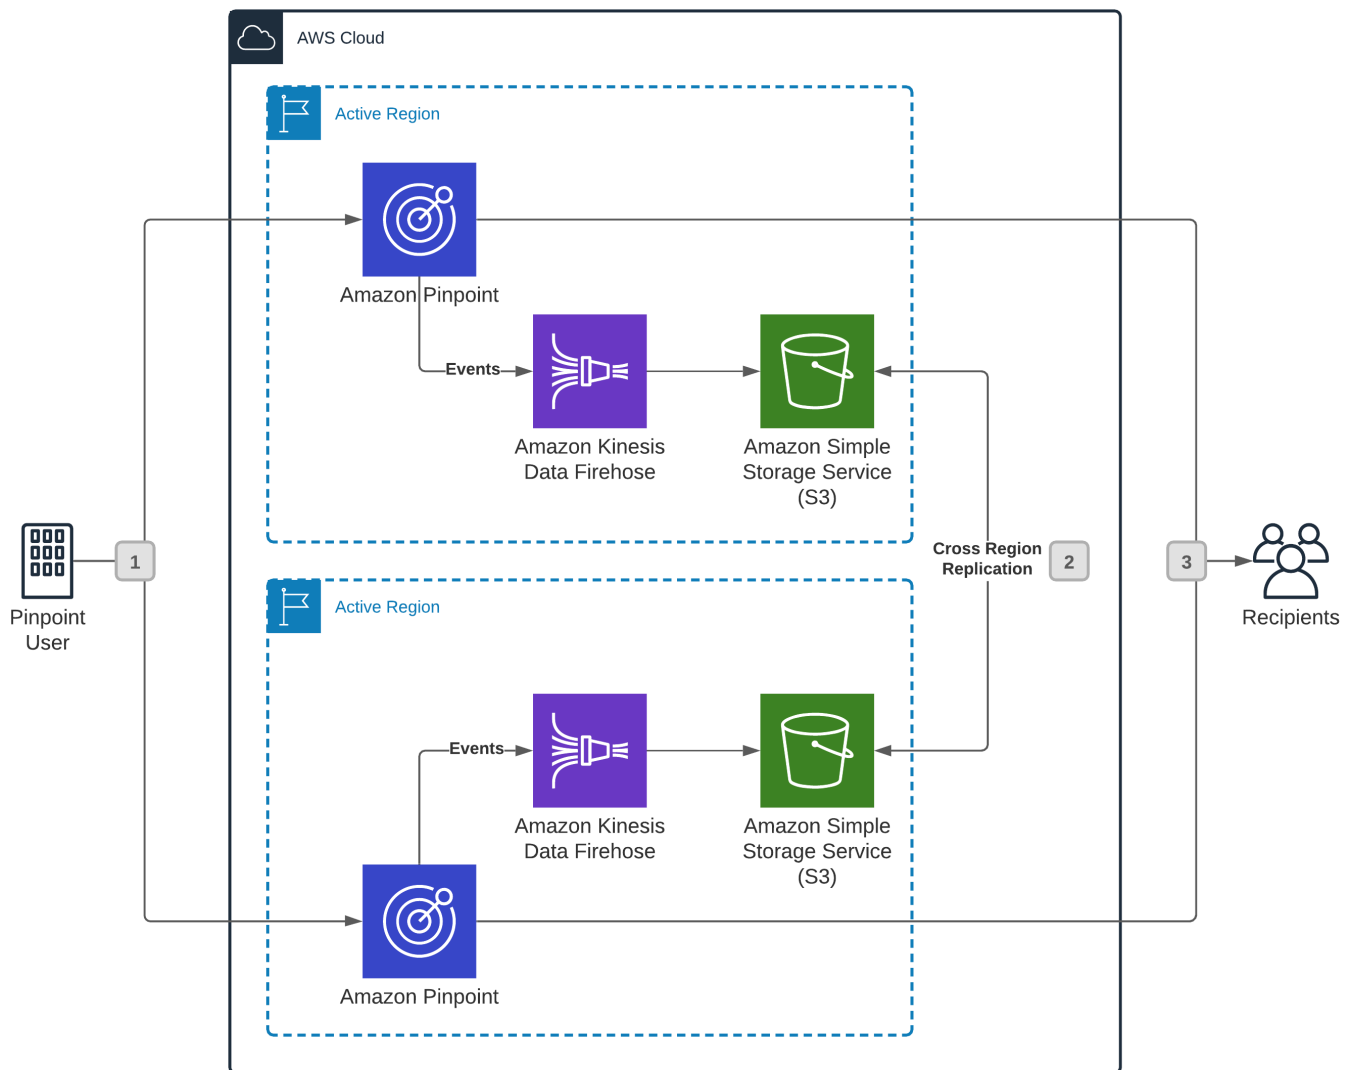
- If you use Amazon Pinpoint to send email, you must configure your sending domains in each AWS Region. This means that you must add multiple records to the DNS configuration of your sending domain.
- If you use Amazon Pinpoint to send SMS messages, you must obtain origination identities in each AWS Region. Origination identities are the identities that are used for sending SMS messages.

Examples of origination identities include short codes, long codes, toll-free numbers, and 10DLC numbers. In some countries, obtaining these origination identities requires a registration process. For example, to obtain a Sender ID for sending SMS messages to recipients in India, you must register your company and use case with regulatory authorities. Similarly, if you want to use a short code to send SMS messages to recipients in the United States, you must register your use case with the mobile carriers.

For more information about requesting various types of origination identities, see [Requesting support for SMS, MMS, and voice messaging](#) in the *Amazon Pinpoint User Guide*. For more information about the types of origination identities that are available in each country, see [Supported countries and regions](#) in the *Amazon Pinpoint User Guide*.

Active-active architecture details

The following diagram shows an active-active architecture for Amazon Pinpoint:



An active-active architecture involves three main parts:

1. Under normal conditions, Amazon Pinpoint traffic is evenly split across two AWS Regions. If there are excessive errors in one Region, then all traffic is routed to the other Region.
2. As messages are sent, Amazon Pinpoint generates [event](#) data. For email messages, this data includes information about deliveries, opens, clicks, bounces, and complaints. For SMS messages, this data includes information about deliveries and failures. This data should be replicated across Regions so that it isn't lost. If you configured your event stream to send data to Amazon S3, you can use Amazon S3 object replication to replicate this data. If your event stream sends data to Amazon Redshift, you can use the cross-region replication feature in Amazon Redshift.

For more information about replicating event data, see [Synchronizing event data](#).

3. Messages are delivered to their recipients.

Benefits and drawbacks of an active-active architecture

Consider the benefits and drawbacks of implementing an active-active architecture.

A potential benefit of this architecture is availability. The Recovery Time Objective (RTO) and Recovery Point Objective (RPO) are minimized in this architecture.

Another benefit of this architecture is that resources are always ready to use. For example, if you have dedicated IP addresses for sending email, those IP addresses are always warmed up because they are constantly being used (assuming that your email is evenly split across both Regions). In a warm standby architecture, the dedicated IPs in the standby Region wouldn't have regular volumes of email being sent from them and would be considered "cold." In a failover scenario, a sudden increase in sending volumes from cold IPs could result in poor deliverability rates.

A drawback of the active-active architecture is cost. Because you're splitting your traffic evenly between two Regions in this architecture, you have to obtain identical resources for both Regions. For example, if you want to use a short code to send SMS messages, you must obtain two separate short codes (one for each Region).

Amazon Pinpoint warm standby reference architecture

This section describes a warm standby architecture for Amazon Pinpoint. In this architecture, a fully functional environment is maintained in an AWS Region that is separate from the primary Region in which you use Amazon Pinpoint. In traditional warm standby architectures, the warm standby Region has reduced capabilities that are scaled up when necessary to meet demand. However, because of the way that Amazon Pinpoint works, it might not be possible to scale up certain resources. For example, resources such as SMS short codes take several weeks to obtain and have a throughput limit that can't be increased on demand. For this reason, you might have to maintain redundant resources in each of the Regions in which you use Amazon Pinpoint.

Topics in this section:

- [Warm standby architecture considerations](#)
- [Warm standby architecture details](#)
- [Benefits and drawbacks of a warm standby architecture](#)

Warm standby architecture considerations

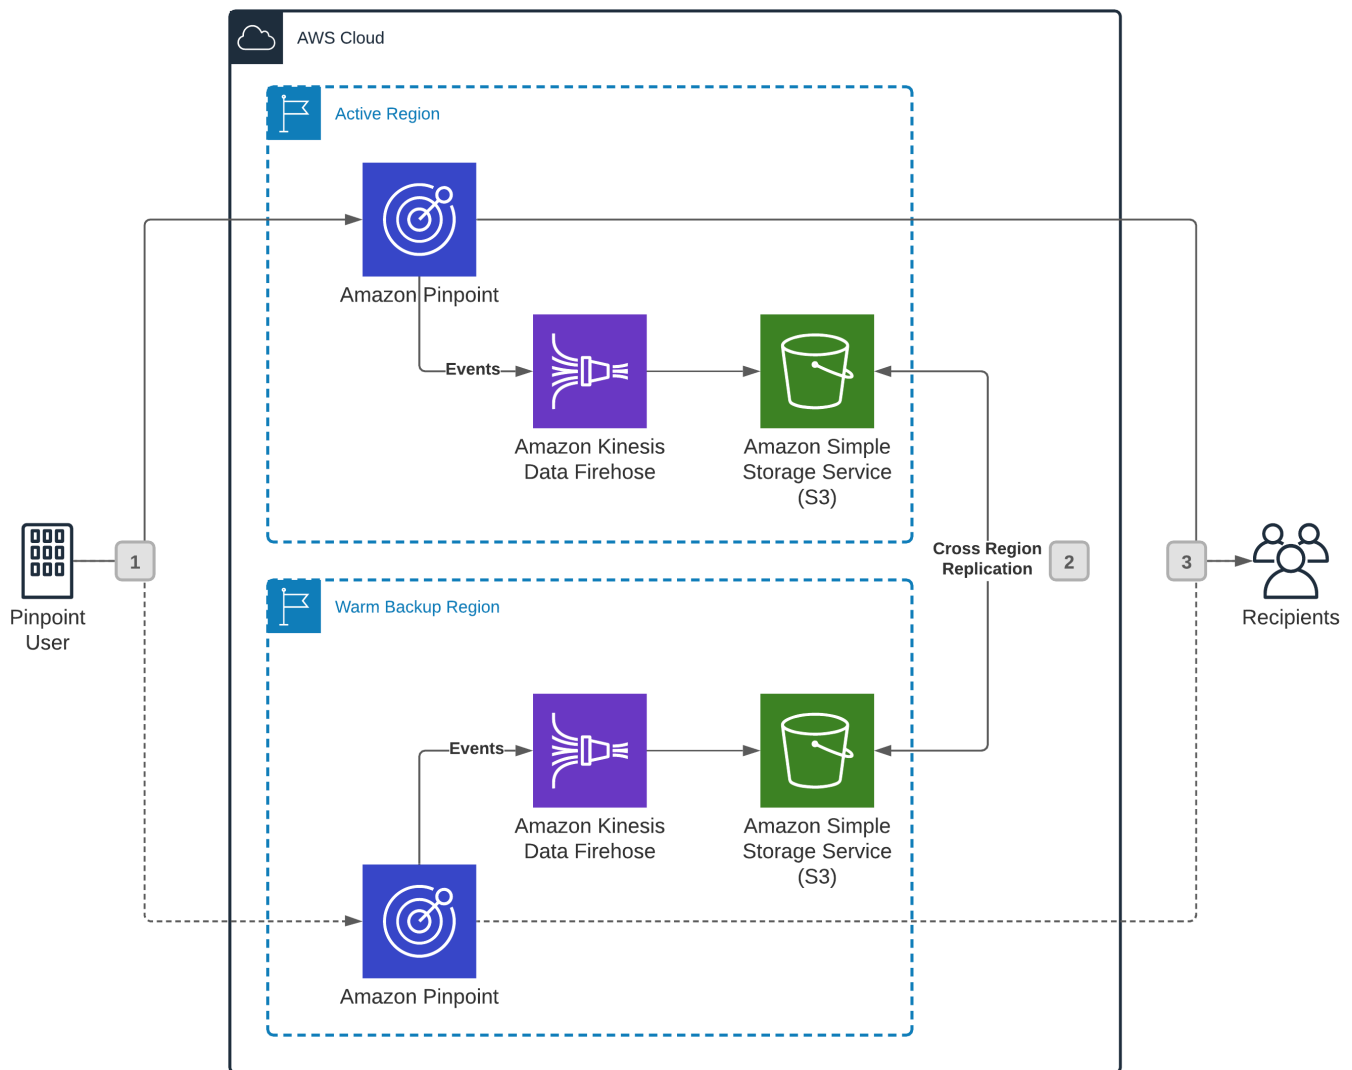
Consider the following factors when implementing a warm standby architecture with Amazon Pinpoint:

- If you use Amazon Pinpoint to send email, you must configure your sending domains in each AWS Region. This means that you must add multiple records to the DNS configuration of your sending domain.
- If you use Amazon Pinpoint to send SMS messages, you must obtain origination identities in each AWS Region. Origination identities are the identities that are used for sending SMS messages. Examples of origination identities include short codes, long codes, toll-free numbers, and 10DLC numbers. In some countries, obtaining these origination identities requires a registration process. For example, to obtain a Sender ID for sending SMS messages to recipients in India, you must register your company and use case with regulatory authorities. Similarly, if you want to use a short code to send SMS messages to recipients in the United States, you must register your use case with the mobile carriers.

For more information about requesting various types of origination identities, see [Requesting SMS support](#) in the *Amazon Pinpoint User Guide*. For more information about the types of origination identities that are available in each country, see [Supported countries and regions](#) in the *Amazon Pinpoint User Guide*.

Warm standby architecture details

The following diagram shows a warm standby architecture for Amazon Pinpoint:



A warm standby architecture involves three main steps:

1. Under normal conditions, Amazon Pinpoint traffic is sent to the primary Region. If there are excessive errors in the primary Region, then all traffic is routed to the other Region.
2. As messages are sent, Amazon Pinpoint generates [event](#) data. For email messages, this data includes information about deliveries, opens, clicks, bounces, and complaints. For SMS messages, this data includes information about deliveries and failures. This data should be replicated across Regions so that it isn't lost. If you configured your event stream to send data to Amazon S3, you can use Amazon S3 object replication to replicate this data. If your event stream sends data to Amazon Redshift, you can use the cross-Region replication feature in Amazon Redshift.

For more information about replicating event data, see [Synchronizing event data](#).

3. Messages are delivered to their recipients.

Benefits and drawbacks of a warm standby architecture

Like any other type of resilient architecture, you must carefully consider the benefits and drawbacks of implementing a warm standby architecture.

A potential benefit of this architecture is cost savings. Because warm standby Regions are only used in rare and temporary situations, it might be possible to provision fewer resources in those Regions. For example, if you use Amazon Pinpoint to send email, you might have dedicated IP addresses (which are available for an additional monthly charge) in your primary Region, but use shared IP addresses (which are available at no additional charge) in your warm standby Region.

A drawback of the warm standby architecture is that you might still need to pay for resources that are unused a majority of the time. For example, if you have a short code for sending SMS messages in your primary Region, and you want your warm standby Region to have the same SMS sending capabilities, then you must provision a short code in the warm standby Region.

Another drawback of the warm standby architecture is that the Recovery Time Objective (RTO) and Recovery Point Objective (RPO) might be slightly higher than they would be for an [active-active](#) architecture. For mission-critical workloads, you should carefully consider both of these architecture options.

Security in Amazon Pinpoint

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from a data center and network architecture that is built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between AWS and you. The [shared responsibility model](#) describes this as security *of* the cloud and security *in* the cloud:

- **Security of the cloud** – AWS is responsible for protecting the infrastructure that runs AWS services in the AWS Cloud. AWS also provides you with services that you can use securely. Third-party auditors regularly test and verify the effectiveness of our security as part of the [AWS Compliance Programs](#). To learn about the compliance programs that apply to Amazon Pinpoint, 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 Amazon Pinpoint. The following topics show you how to configure Amazon Pinpoint to meet your security and compliance objectives. You can also learn how to use other AWS services that help you monitor and secure your Amazon Pinpoint resources.

Topics

- [Data protection in Amazon Pinpoint](#)
- [Identity and access management for Amazon Pinpoint](#)
- [Event logging and monitoring in Amazon Pinpoint](#)
- [Compliance validation for Amazon Pinpoint](#)
- [Infrastructure security in Amazon Pinpoint](#)
- [Security best practices for Amazon Pinpoint](#)

Data protection in Amazon Pinpoint

The AWS [shared responsibility model](#) applies to data protection in Amazon Pinpoint. 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. You are also responsible for the security configuration and management tasks for the AWS services that you use. For more information about data privacy, see the [Data Privacy FAQ](#). For information about data protection in Europe, see the [AWS Shared Responsibility Model and GDPR](#) blog post on the *AWS Security Blog*.

For data protection purposes, we recommend that you protect AWS account credentials and set up individual users with AWS IAM Identity Center or AWS Identity and Access Management (IAM). That way, each user is given only the permissions necessary to fulfill their job duties. We also 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 and recommend TLS 1.3.
- Set up API and user activity logging with AWS CloudTrail. For information about using CloudTrail trails to capture AWS activities, see [Working with CloudTrail trails](#) in the *AWS CloudTrail User Guide*.
- 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 140-3 validated cryptographic modules when accessing AWS through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see [Federal Information Processing Standard \(FIPS\) 140-3](#).

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 Amazon Pinpoint 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.

Depending on how you configure and use the service, Amazon Pinpoint might store the following types of personal data for you or about your customers:

Configuration data

This includes project configuration data such as credentials and settings that define how and when Amazon Pinpoint sends messages through supported channels and the user segments

that it sends messages to. To send messages, this data can include dedicated IP addresses for email messages, short codes and sender IDs for SMS text messages, and credentials for communicating with push notification services such as the Apple Push Notification service (APNs) and Firebase Cloud Messaging (FCM).

User and endpoint data

This includes standard and custom attributes that you use to store and manage data about users and endpoints for an Amazon Pinpoint project. An attribute can store information about a specific user (such as a user's name) or a specific endpoint for a user (such as a user's email address, mobile phone number, or mobile device token). This data can also include external user IDs that correlate users for an Amazon Pinpoint project with users in an external system, such as a customer relationship management system. For more information about what this data can include, see the [User](#) and [Endpoint](#) schemas in the *Amazon Pinpoint API Reference*.

Analytics data

This includes data for metrics, also referred to as *key performance indicators (KPIs)*, that provide insight into the performance of an Amazon Pinpoint project for areas such as user engagement and purchase activity. This also includes data for metrics that provide insight into user demographics for a project. The data can derive from standard and custom attributes for users and endpoints, such as the city where a user lives. It can also derive from events, such as open and click events for the email messages that you send for a project.

Imported data

This includes any user, segmentation, or analytics data that you add or import from external sources and use in Amazon Pinpoint. An example is a JSON file that you import into Amazon Pinpoint (directly through the console or from an Amazon S3 bucket) to build a static segment. Other examples are endpoint data that you add programmatically to build a dynamic segment, endpoint addresses that you send direct messages to, and events that you configure an app to report to Amazon Pinpoint.

Topics

- [Data encryption](#)
- [Internet traffic privacy](#)

Data encryption

Amazon Pinpoint data is encrypted in transit and at rest. When you submit data to Amazon Pinpoint, it encrypts the data as it receives and stores it. When you retrieve data from Amazon Pinpoint, it transmits the data to you by using current security protocols.

Encryption at rest

Amazon Pinpoint encrypts all the data that it stores for you. This includes configuration data, user and endpoint data, analytics data, and any data that you add or import into Amazon Pinpoint. To encrypt your data, Amazon Pinpoint uses internal AWS Key Management Service (AWS KMS) keys that the service owns and maintains on your behalf. We rotate these keys on a regular basis. For information about AWS KMS, see the [AWS Key Management Service Developer Guide](#).

Encryption in transit

Amazon Pinpoint uses HTTPS and Transport Layer Security (TLS) 1.2 or later to communicate with your clients and applications. To communicate with other AWS services, Amazon Pinpoint uses HTTPS and TLS 1.2. In addition, when you create and manage Amazon Pinpoint resources by using the console, an AWS SDK, or the AWS Command Line Interface (AWS CLI), all communications are secured using HTTPS and TLS 1.2.

Key management

To encrypt your Amazon Pinpoint data, Amazon Pinpoint uses internal AWS KMS keys that the service owns and maintains on your behalf. We rotate these keys on a regular basis. You can't provision and use your own AWS KMS or other keys to encrypt data that you store in Amazon Pinpoint.

Internetwork traffic privacy

Internetwork traffic privacy refers to securing connections and traffic between Amazon Pinpoint and your on-premises clients and applications, and between Amazon Pinpoint and other AWS resources in the same AWS Region. The following features and practices can help you achieve internetwork traffic privacy for Amazon Pinpoint.

Traffic between Amazon Pinpoint and on-premises clients and applications

To establish a private connection between Amazon Pinpoint and clients and applications on your on-premises network, you can use Direct Connect. With Direct Connect, you can link your network

to an AWS Direct Connect location by using a standard, fiber-optic Ethernet cable. One end of the cable is connected to your router, and the other end is connected to an Direct Connect router. For more information, see [What is Direct Connect?](#) in the *Direct Connect User Guide*.

To help secure access to Amazon Pinpoint through published API operations, we recommend that you comply with Amazon Pinpoint requirements for API calls. Amazon Pinpoint requires clients to use Transport Layer Security (TLS) 1.2 or later. Clients must also support cipher suites with perfect forward secrecy (PFS), such as Ephemeral Diffie-Hellman (DHE) or Elliptic Curve Diffie-Hellman Ephemeral (ECDHE). Most modern systems such as Java 7 and later, support these modes.

In addition, requests must be signed using an access key ID and a secret access key that's associated with an AWS Identity and Access Management (IAM) principal for your AWS account. Alternatively, you can use the [AWS Security Token Service](#) (AWS STS) to generate temporary security credentials to sign requests.

Traffic between Amazon Pinpoint and other AWS resources

To secure communications between Amazon Pinpoint and other AWS resources in the same AWS Region, Amazon Pinpoint uses HTTPS and TLS 1.2 by default.

Identity and access management for Amazon Pinpoint

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 Amazon Pinpoint resources. IAM is an AWS service that you can use with no additional charge.

Topics

- [Audience](#)
- [Authenticating with identities](#)
- [Managing access using policies](#)
- [How Amazon Pinpoint works with IAM](#)
- [Amazon Pinpoint actions for IAM policies](#)
- [Amazon Pinpoint identity-based policy examples](#)
- [IAM roles for common Amazon Pinpoint tasks](#)
- [Troubleshooting Amazon Pinpoint identity and access management](#)

Audience

How you use AWS Identity and Access Management (IAM) differs based on your role:

- **Service user** - request permissions from your administrator if you cannot access features (see [Troubleshooting Amazon Pinpoint identity and access management](#))
- **Service administrator** - determine user access and submit permission requests (see [How Amazon Pinpoint works with IAM](#))
- **IAM administrator** - write policies to manage access (see [Amazon Pinpoint identity-based policy examples](#))

Authenticating with identities

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

You can sign in as a federated identity using credentials from an identity source like AWS IAM Identity Center (IAM Identity Center), single sign-on authentication, or Google/Facebook credentials. For more information about signing in, see [How to sign in to your AWS account](#) in the *AWS Sign-In User Guide*.

For programmatic access, AWS provides an SDK and CLI to cryptographically sign requests. For more information, see [AWS Signature Version 4 for API requests](#) in the *IAM User Guide*.

AWS account root user

When you create an AWS account, you begin with one sign-in identity called the AWS account *root user* that has complete access to all AWS services and resources. We strongly recommend that you don't use the root user for everyday tasks. For tasks that require root user credentials, see [Tasks that require root user credentials](#) in the *IAM User Guide*.

IAM users and groups

An [IAM user](#) is an identity with specific permissions for a single person or application. We recommend using temporary credentials instead of IAM users with long-term credentials. For more information, see [Require human users to use federation with an identity provider to access AWS using temporary credentials](#) in the *IAM User Guide*.

An [IAM group](#) specifies a collection of IAM users and makes permissions easier to manage for large sets of users. For more information, see [Use cases for IAM users](#) in the *IAM User Guide*.

IAM roles

An [IAM role](#) is an identity with specific permissions that provides temporary credentials. You can assume a role by [switching from a user to an IAM role \(console\)](#) or by calling an AWS CLI or AWS API operation. For more information, see [Methods to assume a role](#) in the *IAM User Guide*.

IAM roles are useful for federated user access, temporary IAM user permissions, cross-account access, cross-service access, and applications running on Amazon EC2. For more information, see [Cross account resource access in IAM](#) in the *IAM User Guide*.

Managing access using policies

You control access in AWS by creating policies and attaching them to AWS identities or resources. A policy defines permissions when associated with an identity or resource. AWS evaluates these policies when a principal makes a request. Most policies are stored in AWS as JSON documents. For more information about JSON policy documents, see [Overview of JSON policies](#) in the *IAM User Guide*.

Using policies, administrators specify who has access to what by defining which **principal** can perform **actions** on what **resources**, and under what **conditions**.

By default, users and roles have no permissions. An IAM administrator creates IAM policies and adds them to roles, which users can then assume. IAM policies define permissions regardless of the method used to perform the operation.

Identity-based policies

Identity-based policies are JSON permissions policy documents that you attach to an identity (user, group, or role). These policies control what actions identities can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see [Define custom IAM permissions with customer managed policies](#) in the *IAM User Guide*.

Identity-based policies can be *inline policies* (embedded directly into a single identity) or *managed policies* (standalone policies attached to multiple identities). To learn how to choose between managed and inline policies, see [Choose between managed policies and inline policies](#) in the *IAM User Guide*.

Amazon Pinpoint supports the use of identity-based policies to control access to Amazon Pinpoint resources.

Resource-based policies

Resource-based policies are JSON policy documents that you attach to a resource. Examples include IAM *role trust policies* and Amazon S3 *bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. You must [specify a principal](#) in a resource-based policy.

Resource-based policies are inline policies that are located in that service. You can't use AWS managed policies from IAM in a resource-based policy.

Amazon Pinpoint supports the use of resource-based policies to control access to Amazon Pinpoint resources.

Access control lists (ACLs)

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

Amazon S3, AWS WAF, and Amazon VPC are examples of services that support ACLs. To learn more about ACLs, see [Access control list \(ACL\) overview](#) in the *Amazon Simple Storage Service Developer Guide*.

Amazon Pinpoint doesn't support the use of ACLs to control access to Amazon Pinpoint resources.

Other policy types

AWS supports additional policy types that can set the maximum permissions granted by more common policy types:

- **Permissions boundaries** – Set the maximum permissions that an identity-based policy can grant to an IAM entity. For more information, see [Permissions boundaries for IAM entities](#) in the *IAM User Guide*.
- **Service control policies (SCPs)** – Specify the maximum permissions for an organization or organizational unit in AWS Organizations. For more information, see [Service control policies](#) in the *AWS Organizations User Guide*.

- **Resource control policies (RCPs)** – Set the maximum available permissions for resources in your accounts. For more information, see [Resource control policies \(RCPs\)](#) in the *AWS Organizations User Guide*.
- **Session policies** – Advanced policies passed as a parameter when creating a temporary session for a role or federated user. For more information, see [Session policies](#) in the *IAM User Guide*.

Amazon Pinpoint supports the use of these types of policies to control access to Amazon Pinpoint resources.

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 Pinpoint works with IAM

To use Amazon Pinpoint, users in your AWS account require permissions that allow them to view analytics data, create projects, define user segments, deploy campaigns, and more. If you integrate a mobile or web app with Amazon Pinpoint, users of your app also require access to Amazon Pinpoint. This access enables your app to register endpoints and report usage data to Amazon Pinpoint. To grant access to Amazon Pinpoint features, create AWS Identity and Access Management (IAM) policies that allow Amazon Pinpoint actions for IAM identities or Amazon Pinpoint resources.

IAM is a service that helps administrators securely control access to AWS resources. IAM policies include statements that allow or deny specific actions by specific users or for specific resources. Amazon Pinpoint provides a [set of actions](#) that you can use in IAM policies to specify specific permissions for Amazon Pinpoint users and resources. This means that you can grant the appropriate level of access to Amazon Pinpoint without creating overly permissive policies that might expose important data or compromise your resources. For example, you can grant unrestricted access to an Amazon Pinpoint administrator, and grant read-only access to individuals who only need access to a specific project.

Before you use IAM to manage access to Amazon Pinpoint, you should understand what IAM features are available for use with Amazon Pinpoint. To get a high-level view of how Amazon Pinpoint and other AWS services work with IAM, see [AWS services that work with IAM](#) in the *IAM User Guide*.

Topics

- [Amazon Pinpoint identity-based policies](#)
- [Amazon Pinpoint resource-based policies](#)
- [Authorization based on Amazon Pinpoint tags](#)
- [Amazon Pinpoint IAM roles](#)

Amazon Pinpoint identity-based policies

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. Amazon Pinpoint supports specific actions, resources, and condition keys. To learn about all the elements that you can use in a JSON policy, see [IAM JSON policy elements reference](#) in the *IAM User Guide*.

Actions

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Action` element of a JSON policy describes the actions that you can use to allow or deny access in a policy. Include actions in a policy to grant permissions to perform the associated operation.

This means that policy actions control what users can do on the Amazon Pinpoint console. Policy actions also control what users can do programmatically by directly using the AWS SDKs, the AWS Command Line Interface (AWS CLI), or the Amazon Pinpoint APIs.

Policy actions in Amazon Pinpoint use the following prefixes:

- **mobiletargeting** – For actions that derive from the Amazon Pinpoint API, which is the primary API for Amazon Pinpoint.
- **sms-voice** – For actions that derive from the Amazon Pinpoint SMS and Voice API, which is a supplemental API that provides advanced options for using and managing the SMS and voice channels in Amazon Pinpoint.

For example, to grant someone permission to view information about all the segments for a project, which is an action that corresponds to the `GetSegments` operation in the Amazon Pinpoint API, include the `mobiletargeting:GetSegments` action in their policy. Policy

statements must include either an `Action` or `NotAction` element. Amazon Pinpoint defines its own set of actions that describe the tasks that users can perform with it.

To specify multiple actions in a single statement, separate them with commas:

```
"Action": [
  "mobiletargeting:action1",
  "mobiletargeting:action2"
```

You can also specify multiple actions by using wildcards (*). For example, to specify all actions that begin with the word `Get`, include the following action:

```
"Action": "mobiletargeting:Get*"
```

However, as a best practice, you should create policies that follow the principle of *least privilege*. In other words, you should create policies that include only the permissions that are required to perform a specific action.

For a list of Amazon Pinpoint actions that you can use in IAM policies, see [Amazon Pinpoint actions for IAM policies](#).

Resources

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Resource` JSON policy element specifies the object or objects to which the action applies. As a best practice, specify a resource using its [Amazon Resource Name \(ARN\)](#). For actions that don't support resource-level permissions, use a wildcard (*) to indicate that the statement applies to all resources.

```
"Resource": "*" 
```

For example, the `mobiletargeting:GetSegments` action retrieves information about all the segments that are associated with a specific Amazon Pinpoint project. You identify a project with an ARN in the following format:

```
arn:aws:mobiletargeting:${Region}:${Account}:apps/${projectId}
```

For more information about the format of ARNs, see [Amazon Resource Names \(ARNs\)](#) in the *AWS General Reference*.

In IAM policies, you can specify ARNs for the following types of Amazon Pinpoint resources:

- Campaigns
- Journeys
- Message templates (referred to as *templates* in some contexts)
- Projects (referred to as *apps* or *applications* in some contexts)
- Recommender models (referred to as *recommenders* in some contexts)
- Segments

For example, to create a policy statement for the project that has the project ID `810c7aab86d42fb2b56c8c966example`, use the following ARN:

```
"Resource": "arn:aws:mobiletargeting:us-east-1:123456789012:apps/810c7aab86d42fb2b56c8c966example"
```

To specify all the projects that belong to a specific account, use the wildcard (*):

```
"Resource": "arn:aws:mobiletargeting:us-east-1:123456789012:apps/*"
```

Some Amazon Pinpoint actions, such as certain actions for creating resources, can't be performed on a specific resource. In those cases, you must use the wildcard (*):

```
"Resource": "*"
```

In IAM policies, you can also specify ARNs for the following types of Amazon Pinpoint SMS and Voice resources:

- Configuration Set
- Opt Out List
- Phone Number
- Pool
- Sender Id

For example, to create a policy statement for a phone number that has the phone number ID phone-12345678901234567890123456789012 use the following ARN:

```
"Resource": "arn:aws:sms-voice:us-east-1:123456789012:phone-number/  
phone-12345678901234567890123456789012"
```

To specify all phone numbers that belong to a specific account, use a wildcard (*) in place of the phone number ID:

```
"Resource": "arn:aws:sms-voice:us-east-1:123456789012:phone-number/*"
```

Some Amazon Pinpoint SMS and Voice actions are not performed on a specific resource, such as those for managing account-level settings like spend limits. In those cases, you must use the wildcard (*):

```
"Resource": "*" 
```

Some Amazon Pinpoint API actions involve multiple resources. For example, the `TagResource` action can add a tag to multiple projects. To specify multiple resources in a single statement, separate the ARNs with commas:

```
"Resource": [  
    "resource1",  
    "resource2"
```

To see a list of Amazon Pinpoint resource types and their ARNs, see [Resources Defined by Amazon Pinpoint](#) in the *IAM User Guide*. To learn which actions you can specify with the ARN of each resource type, see [Actions Defined by Amazon Pinpoint](#) in the *IAM User Guide*.

Condition keys

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Condition` element specifies when statements execute based on defined criteria. You can create conditional expressions that use [condition operators](#), such as equals or less than, to match the condition in the policy with values in the request. To see all AWS global condition keys, see [AWS global condition context keys](#) in the *IAM User Guide*.

Amazon Pinpoint defines its own set of condition keys and also supports some global condition keys. To see a list of all AWS global condition keys, see [AWS global condition context keys](#) in the *IAM User Guide*. To see a list of Amazon Pinpoint condition keys, see [Condition Keys for Amazon Pinpoint](#) in the *IAM User Guide*. To learn which actions and resources you can use a condition key with, see [Actions Defined by Amazon Pinpoint](#) in the *IAM User Guide*.

Examples

To view examples of Amazon Pinpoint identity-based policies, see [Amazon Pinpoint identity-based policy examples](#).

Amazon Pinpoint resource-based policies

Resource-based policies are JSON policy documents that specify what actions a specified principal can perform on an Amazon Pinpoint resource and under what conditions. Amazon Pinpoint supports resource-based permissions policies for campaigns, journeys, message templates (*templates*), recommender models (*recommenders*), projects (*apps*), and segments. Resource-based policies let you grant usage permission to other accounts on a per-resource basis. You can also use a resource-based policy to allow another AWS service to access these types of Amazon Pinpoint resources.

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, you must also grant the principal entity permission to access the resource. 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 [Cross account resource access in IAM](#) in the *IAM User Guide*.

Examples

To view examples of Amazon Pinpoint resource-based policies, see [the section called “Identity-based policy examples”](#).

Authorization based on Amazon Pinpoint tags

You can associate tags with certain types of Amazon Pinpoint resources or pass tags in a request to Amazon Pinpoint. To control access based on tags, you provide tag information in the [condition](#)

[element](#) of a policy using the `aws:ResourceTag/${TagKey}`, `aws:RequestTag/${TagKey}`, or `aws:TagKeys` condition keys.

For information about tagging Amazon Pinpoint resources, including an example IAM policy, see [Tagging resources](#) in the *Amazon Pinpoint Developer Guide*.

Amazon Pinpoint IAM roles

An [IAM role](#) is an entity within your AWS account that has specific permissions.

Using temporary credentials with Amazon Pinpoint

You can use temporary credentials to sign in with federation, assume an IAM role, or assume a cross-account role. You obtain temporary security credentials by calling AWS Security Token Service (AWS STS) API operations such as [AssumeRole](#) or [GetFederationToken](#).

Amazon Pinpoint supports using temporary credentials.

Service-linked roles

[Service-linked roles \(SLR\)](#) allow AWS services to access resources in other services to complete an action on your behalf. Service-linked roles appear in your IAM account and are owned by the service. An IAM administrator can view but not edit the permissions for service-linked roles.

Amazon Pinpoint doesn't use service-linked roles.

Service roles

This feature allows a service to assume a [service role](#) on your behalf. This role allows the service to access resources in other services to complete an action on your behalf. Service roles appear in your IAM account and are owned by the account. This means that an IAM administrator can change the permissions for this role. However, doing so might break the functionality of the service.

Amazon Pinpoint supports using service roles.

Amazon Pinpoint actions for IAM policies

To manage access to Amazon Pinpoint resources in your AWS account, you can add Amazon Pinpoint actions to AWS Identity and Access Management (IAM) policies. By using actions in policies, you can control what users can do on the Amazon Pinpoint console. You can also control

what users can do programmatically by using the AWS SDKs, the AWS Command Line Interface (AWS CLI), or the Amazon Pinpoint API operations directly.

In a policy, you specify each action with the appropriate Amazon Pinpoint namespace, followed by a colon (:) and the name of the action, such as `GetSegments`. Most actions correspond to a request to the Amazon Pinpoint API using a specific URI and HTTP method. For example, if you allow the `mobiletargeting:GetSegments` action in a user's policy, the user is allowed to retrieve information about all the segments for a project by submitting an HTTP GET request to the [apps/projectId/segments](#) URI. This policy also allows the user to view that information on the console, and retrieve that information by using an AWS SDK or the AWS CLI.

Each action is performed on a specific Amazon Pinpoint resource, which you identify in a policy statement by its Amazon Resource Name (ARN). For example, the `mobiletargeting:GetSegments` action is performed on a specific project, which you identify with the ARN, `arn:aws:mobiletargeting:region:accountId:apps/projectId`.

This topic identifies Amazon Pinpoint actions that you can add to IAM policies for your AWS account. For examples that demonstrate how you can use actions in policies to manage access to Amazon Pinpoint resources, see [Amazon Pinpoint identity-based policy examples](#).

Topics

- [Amazon Pinpoint API actions](#)
- [Amazon Pinpoint SMS and voice version 1 API actions](#)
- [AWS End User Messaging SMS and voice version 2 API actions](#)

Amazon Pinpoint API actions

This section identifies actions for features that are available from the Amazon Pinpoint API, which is the primary API for Amazon Pinpoint. To learn more about this API, see the [Amazon Pinpoint API Reference](#).

Categories:

- [Analytics and metrics](#)
- [Campaigns](#)
- [Channels](#)
- [Endpoints](#)
- [Event streams](#)

- [Events](#)
- [Export jobs](#)
- [Import jobs](#)
- [Journeys](#)
- [Message templates](#)
- [Messages](#)
- [One-time passwords](#)
- [Phone number validation](#)
- [Projects](#)
- [Recommender models](#)
- [Segments](#)
- [Tags](#)
- [Users](#)

Analytics and metrics

The following permissions are related to viewing analytics data on the Amazon Pinpoint console. They're also related to retrieving (querying) aggregated data for standard metrics, also referred to as *key performance indicators (KPIs)*, that apply to projects, campaigns, and journeys.

mobiletargeting:GetReports

View analytics data on the Amazon Pinpoint console. This permission is also required in order to create segments that contain custom attributes using the Amazon Pinpoint console. It's also required to obtain an estimate of the size of a segment in the Amazon Pinpoint console.

- URI – Not applicable
- Method – Not applicable
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:*`

mobiletargeting:GetApplicationDateRangeKpi

Retrieve (query) aggregated data for a standard application metric. This is a metric that applies to all the campaigns or transactional messages that are associated with a project.

- URI – `/apps/projectId/kpis/daterange/kpi-name`
- Method – GET

- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/kpis/daterange/kpi-name`

mobiletargeting:GetCampaignDateRangeKpi

Retrieve (query) aggregated data for a standard campaign metric. This is a metric that applies to an individual campaign.

- URI – [/apps/projectId/campaigns/campaignId/kpis/daterange/kpi-name](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/campaigns/campaignId/kpis/daterange/kpi-name`

mobiletargeting:GetJourneyDateRangeKpi

Retrieve (query) aggregated data for a standard journey engagement metric. This is an engagement metric that applies to an individual journey—for example, the number of messages that were opened by participants for all the activities in a journey.

- URI – [/apps/projectId/journeys/journeyId/kpis/daterange/kpi-name](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/journeys/journeyId/kpis/daterange/kpi-name`

mobiletargeting:GetJourneyExecutionMetrics

Retrieve (query) aggregated data for standard execution metrics that apply to an individual journey—for example, the number of participants who are actively proceeding through all the activities in a journey.

- URI – [/apps/projectId/journeys/journeyId/execution-metrics](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/journeys/journeyId/execution-metrics`

mobiletargeting:GetJourneyExecutionActivityMetrics

Retrieve (query) aggregated data for standard execution metrics that apply to an individual activity in a journey—for example, the number of participants who started or completed an activity.

- URI – [/apps/projectId/journeys/journeyId/activities/journey-activity-id/execution-metrics](#)

- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/journeys/journeyId/activities/journey-activity-id/execution-metrics`

Campaigns

The following permissions are related to managing campaigns in your Amazon Pinpoint account.

mobiletargeting:CreateCampaign

Create a campaign for a project.

- URI – [/apps/projectId/campaigns](#)
- Method – POST
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/campaigns`

mobiletargeting>DeleteCampaign

Delete a specific campaign.

- URI – [/apps/projectId/campaigns/campaignId](#)
- Method – DELETE
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/campaigns/campaignId`

mobiletargeting:GetCampaign

Retrieve information about a specific campaign.

- URI – [/apps/projectId/campaigns/campaignId](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/campaigns/campaignId`

mobiletargeting:GetCampaignActivities

Retrieve information about the activities performed by a campaign.

- URI – [/apps/projectId/campaigns/campaignId/activities](#)
- Method – GET

- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/campaigns/campaignId`

mobiletargeting:GetCampaigns

Retrieve information about all campaigns for a project.

- URI – [/apps/projectId/campaigns](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/campaigns/campaignId`

mobiletargeting:GetCampaignVersion

Retrieve information about a specific campaign version.

- URI – [/apps/projectId/campaigns/campaignId/versions/versionId](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/campaigns/campaignId`

mobiletargeting:GetCampaignVersions

Retrieve information about the current and prior versions of a campaign.

- URI – [/apps/projectId/campaigns/campaignId/versions](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/campaigns/campaignId`

mobiletargeting:UpdateCampaign

Update a specific campaign.

- URI – [/apps/projectId/campaigns/campaignId](#)
- Method – PUT
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/campaigns/campaignId`

Channels

The following permissions are related to managing channels in your Amazon Pinpoint account. In Amazon Pinpoint, *channels* refer to the methods that you use to contact your customers, such as sending email, SMS messages, or push notifications.

mobiletargeting:DeleteAdmChannel

Deactivate the Amazon Device Messaging (ADM) channel for a project.

- URI – [/apps/projectId/channels/adm](#)
- Method – DELETE
- Resource ARN – arn:aws:mobiletargeting:region:accountId:apps/projectId/channels/adm

mobiletargeting:GetAdmChannel

Retrieve information about the ADM channel for a project.

- URI – [/apps/projectId/channels/adm](#)
- Method – GET
- Resource ARN – arn:aws:mobiletargeting:region:accountId:apps/projectId/channels/adm

mobiletargeting:UpdateAdmChannel

Activate or update the ADM channel for a project.

- URI – [/apps/projectId/channels/adm](#)
- Method – PUT
- Resource ARN – arn:aws:mobiletargeting:region:accountId:apps/projectId/channels/adm

mobiletargeting>DeleteApnsChannel

Deactivate the Apple Push Notification service (APNs) channel for a project.

- URI – [/apps/projectId/channels/apns](#)
- Method – DELETE
- Resource ARN – arn:aws:mobiletargeting:region:accountId:apps/projectId/channels/apns

mobiletargeting:GetApnsChannel

Retrieve information about the APNs channel for a project.

- URI – [/apps/projectId/channels/apns](#)
- Method – GET

- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/channels/apns`

mobiletargeting:UpdateApnsChannel

Activate or update the APNs channel for a project.

- URI – [/apps/projectId/channels/apns](#)
- Method – PUT
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/channels/apns`

mobiletargeting>DeleteApnsSandboxChannel

Deactivate the APNs sandbox channel for a project.

- URI – [/apps/projectId/channels/apns_sandbox](#)
- Method – DELETE
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/channels/apns_sandbox`

mobiletargeting:GetApnsSandboxChannel

Retrieve information about the APNs sandbox channel for a project.

- URI – [/apps/projectId/channels/apns_sandbox](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/channels/apns_sandbox`

mobiletargeting:UpdateApnsSandboxChannel

Activate or update the APNs sandbox channel for a project.

- URI – [/apps/projectId/channels/apns_sandbox](#)
- Method – PUT
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/channels/apns_sandbox`

mobiletargeting>DeleteApnsVoipChannel

Deactivate the APNs VoIP channel for a project.

- URI – [/apps/projectId/channels/apns_voip](#)

- Method – DELETE
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/channels/apns_voip`

mobiletargeting:GetApnsVoipChannel

Retrieve information about the APNs VoIP channel for a project.

- URI – [/apps/projectId/channels/apns_voip](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/channels/apns_voip`

mobiletargeting:UpdateApnsVoipChannel

Activate or update the APNs VoIP channel for a project.

- URI – [/apps/projectId/channels/apns_voip](#)
- Method – PUT
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/channels/apns_voip`

mobiletargeting>DeleteApnsVoipSandboxChannel

Deactivate the APNs VoIP sandbox channel for a project.

- URI – [/apps/projectId/channels/apns_voip_sandbox](#)
- Method – DELETE
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/channels/apns_voip_sandbox`

mobiletargeting:GetApnsVoipSandboxChannel

Retrieve information about the APNs VoIP sandbox channel for a project.

- URI – [/apps/projectId/channels/apns_voip_sandbox](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/channels/apns_voip_sandbox`

mobiletargeting:UpdateApnsVoipSandboxChannel

Activate or update the APNs VoIP sandbox channel for a project.

- URI – [/apps/*projectId*/channels/apns_voip_sandbox](#)
- Method – PUT
- Resource ARN – arn:aws:mobiletargeting:*region*:*accountId*:apps/*projectId*/channels/apns_voip_sandbox

mobiletargeting:DeleteBaiduChannel

Deactivate the Baidu Cloud Push channel for a project.

- URI – [/apps/*projectId*/channels/baidu](#)
- Method – DELETE
- Resource ARN – arn:aws:mobiletargeting:*region*:*accountId*:apps/*projectId*/channels/baidu

mobiletargeting:GetBaiduChannel

Retrieve information about the Baidu Cloud Push channel for a project.

- URI – [/apps/*projectId*/channels/baidu](#)
- Method – GET
- Resource ARN – arn:aws:mobiletargeting:*region*:*accountId*:apps/*projectId*/channels/baidu

mobiletargeting:UpdateBaiduChannel

Activate or update the Baidu Cloud Push channel for a project.

- URI – [/apps/*projectId*/channels/baidu](#)
- Method – PUT
- Resource ARN – arn:aws:mobiletargeting:*region*:*accountId*:apps/*projectId*/channels/baidu

mobiletargeting>DeleteEmailChannel

Deactivate the email channel for a project.

- URI – [/apps/*projectId*/channels/email](#)
- Method – DELETE
- Resource ARN – arn:aws:mobiletargeting:*region*:*accountId*:apps/*projectId*/channels/email

mobiletargeting:GetEmailChannel

Retrieve information about the email channel for a project.

- URI – [/apps/*projectId*/channels/email](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/channels/email`

mobiletargeting:UpdateEmailChannel

Activate or update the email channel for a project.

- URI – [/apps/*projectId*/channels/email](#)
- Method – PUT
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/channels/email`

mobiletargeting>DeleteGcmChannel

Deactivate the Firebase Cloud Messaging (FCM) channel for a project. This channel allows Amazon Pinpoint to send push notifications to an Android app through the FCM service, which replaces the Google Cloud Messaging (GCM) service.

- URI – [/apps/*projectId*/channels/gcm](#)
- Method – DELETE
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/channels/gcm`

mobiletargeting:GetGcmChannel

Retrieve information about the FCM channel for a project. This channel allows Amazon Pinpoint to send push notifications to an Android app through the FCM service, which replaces the Google Cloud Messaging (GCM) service.

- URI – [/apps/*projectId*/channels/gcm](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/channels/gcm`

mobiletargeting:UpdateGcmChannel

Activate or update the FCM channel for a project. This channel allows Amazon Pinpoint to send push notifications to an Android app through the FCM service, which replaces the Google Cloud Messaging (GCM) service.

- URI – [/apps/*projectId*/channels/gcm](#)
- Method – PUT
- Resource ARN – arn:aws:mobiletargeting:*region*:*accountId*:apps/*projectId*/channels/gcm

mobiletargeting:DeleteSmsChannel

Deactivate the SMS channel for a project.

- URI – [/apps/*projectId*/channels/sms](#)
- Method – DELETE
- Resource ARN – arn:aws:mobiletargeting:*region*:*accountId*:apps/*projectId*/channels/sms

mobiletargeting:GetSmsChannel

Retrieve information about the SMS channel for a project.

- URI – [/apps/*projectId*/channels/sms](#)
- Method – GET
- Resource ARN – arn:aws:mobiletargeting:*region*:*accountId*:apps/*projectId*/channels/sms

mobiletargeting:UpdateSmsChannel

Activate or update the SMS channel for a project.

- URI – [/apps/*projectId*/channels/sms](#)
- Method – PUT
- Resource ARN – arn:aws:mobiletargeting:*region*:*accountId*:apps/*projectId*/channels/sms

mobiletargeting:GetChannels

Retrieves information about the history and status of each channel for an application.

- URI – [/apps/*application-id*/channels](#)
- Method – GET
- Resource ARN – arn:aws:mobiletargeting:*region*:*accountId*:apps/*projectId*/channels

mobiletargeting:DeleteVoiceChannel

Deactivate the voice channel for an application and deletes any existing settings for the channel.

- URI – [/apps/application-id/channels/voice](#)
- Method – DELETE
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectid/channels/voice`

mobiletargeting:GetVoiceChannel

Retrieves information about the status and settings of the voice channel for an application.

- URI – [/apps/application-id/channels/voice](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectid/channels/voice`

mobiletargeting:UpdateVoiceChannel

Enables the voice channel for an application or updates the status and settings of the voice channel for an application.

- URI – [/apps/application-id/channels/voice](#)
- Method – PUT
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectid/channels/voice`

Endpoints

The following permissions are related to managing endpoints in your Amazon Pinpoint account. In Amazon Pinpoint, an *endpoint* is a single destination for your messages. For example, an endpoint could be a customer's email address, telephone number, or mobile device token.

mobiletargeting>DeleteEndpoint

Delete an endpoint.

- URI – [/apps/projectId/endpoints/endpointId](#)
- Method – DELETE

- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/endpoints/endpointId`

mobiletargeting:GetEndpoint

Retrieve information about a specific endpoint.

- URI – [/apps/projectId/endpoints/endpointId](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/endpoints/endpointId`

mobiletargeting:RemoveAttributes

Remove one or more attributes, of the same attribute type, from all the endpoints that are associated with an application.

- URI – [apps/application-id/attributes/attribute-type](#)
- Method – PUT
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/attributes/attribute-type`

mobiletargeting:UpdateEndpoint

Create an endpoint or update the information for an endpoint.

- URI – [/apps/projectId/endpoints/endpointId](#)
- Method – PUT
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/endpoints/endpointId`

mobiletargeting:UpdateEndpointsBatch

Create or update endpoints as a batch operation.

- URI – [/apps/projectId/endpoints](#)
- Method – PUT
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId`

Event streams

The following permissions are related to managing event streams for your Amazon Pinpoint account.

mobiletargeting:DeleteEventStream

Delete the event stream for a project.

- URI – [/apps/*projectId*/eventstream/](#)
- Method – DELETE
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/eventstream`

mobiletargeting:GetEventStream

Retrieve information about the event stream for a project.

- URI – [/apps/*projectId*/eventstream/](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/eventstream`

mobiletargeting:PutEventStream

Create or update an event stream for a project.

- URI – [/apps/*projectId*/eventstream/](#)
- Method – POST
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/eventstream`

Events

The following permissions are related to managing events jobs in your Amazon Pinpoint account. In Amazon Pinpoint, you create *import jobs* to create segments based on endpoint definitions that are stored in an Amazon S3 bucket.

mobiletargeting:PutEvents

Creates a new event to record for endpoints, or creates or updates endpoint data that existing events are associated with.

- URI – [/apps/*application-id*/events](#)
- Method – POST
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/events`

Export jobs

The following permissions are related to managing export jobs in your Amazon Pinpoint account. In Amazon Pinpoint, you create *export jobs* to send information about endpoints to an Amazon S3 bucket for storage or analysis.

mobiletargeting:CreateExportJob

Create an export job for exporting endpoint definitions to Amazon S3.

- URI – [/apps/*projectId*/jobs/export](#)
- Method – POST
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/jobs/export`

mobiletargeting:GetExportJob

Retrieve information about a specific export job for a project.

- URI – [/apps/*projectId*/jobs/export/*jobId*](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/jobs/export/jobId`

mobiletargeting:GetExportJobs

Retrieve a list of all the export jobs for a project.

- URI – [/apps/*projectId*/jobs/export](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/jobs/export`

Import jobs

The following permissions are related to managing import jobs in your Amazon Pinpoint account. In Amazon Pinpoint, you create *import jobs* to create segments based on endpoint definitions that are stored in an Amazon S3 bucket.

mobiletargeting:CreateImportJob

Import endpoint definitions from Amazon S3 to create a segment.

- URI – [/apps/*projectId*/jobs/import](#)
- Method – POST
- Resource ARN – arn:aws:mobiletargeting:*region*:*accountId*:apps/*projectId*

mobiletargeting:GetImportJob

Retrieve information about a specific import job for a project.

- URI – [/apps/*projectId*/jobs/import/*jobId*](#)
- Method – GET
- Resource ARN – arn:aws:mobiletargeting:*region*:*accountId*:apps/*projectId*/jobs/import/*jobId*

mobiletargeting:GetImportJobs

Retrieve information about all the import jobs for a project.

- URI – [/apps/*projectId*/jobs/import](#)
- Method – GET
- Resource ARN – arn:aws:mobiletargeting:*region*:*accountId*:apps/*projectId*

Journeys

The following permissions are related to managing journeys in your Amazon Pinpoint account.

mobiletargeting:CreateJourney

Create a journey for a project.

- URI – [/apps/*projectId*/journeys](#)
- Method – POST
- Resource ARN – arn:aws:mobiletargeting:*region*:*accountId*:apps/*projectId*/journeys

mobiletargeting:GetJourney

Retrieve information about a specific journey.

- URI – [/apps/*projectId*/journeys/*journeyId*](#)
- Method – GET
- Resource ARN – arn:aws:mobiletargeting:*region*:*accountId*:apps/*projectId*/journeys/*journeyId*

mobiletargeting:ListJourneys

Retrieve information about all the journeys for a project.

- URI – [/apps/projectId/journeys](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/journeys`

mobiletargeting:UpdateJourney

Update the configuration and other settings for a specific journey.

- URI – [/apps/projectId/journeys/journeyId](#)
- Method – PUT
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/journeys/journeyId`

mobiletargeting:UpdateJourneyState

Cancel an active journey.

- URI – [/apps/projectId/journeys/journeyId/state](#)
- Method – PUT
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/journeys/journeyId/state`

mobiletargeting>DeleteJourney

Delete a specific journey.

- URI – [/apps/projectId/journeys/journeyId](#)
- Method – DELETE
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/journeys/journeyId`

Message templates

The following permissions are related to creating and managing message templates for your Amazon Pinpoint account. A *message template* is a set of content and settings that you can define, save, and reuse in messages that you send for any of your Amazon Pinpoint projects.

mobiletargeting:ListTemplates

Retrieve information about all the message templates that are associated with your Amazon Pinpoint account.

- URI – [/templates](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:templates`

mobiletargeting:ListTemplateVersions

Retrieve information about all the versions of a specific message template.

- URI – [/templates/template-name/template-type/versions](#)
- Method – GET
- Resource ARN – Not applicable

mobiletargeting:UpdateTemplateActiveVersion

Designate a specific version of a message template as the active version of the template.

- URI – [/templates/template-name/template-type/active-version](#)
- Method – GET
- Resource ARN – Not applicable

mobiletargeting:GetEmailTemplate

Retrieve information about a message template for messages that are sent through the email channel.

- URI – [/templates/template-name/email](#)
- Method – GET
- Resource ARN –
`arn:aws:mobiletargeting:region:accountId:templates/template-name/EMAIL`

mobiletargeting>CreateEmailTemplate

Create a message template for messages that are sent through the email channel.

- URI – [/templates/*template-name*/email](#)
- Method – POST
- Resource ARN –
arn:aws:mobiletargeting:*region*:*accountId*:templates/*template-name*/EMAIL

mobiletargeting:UpdateEmailTemplate

Update an existing message template for messages that are sent through the email channel.

- URI – [/templates/*template-name*/email](#)
- Method – PUT
- Resource ARN –
arn:aws:mobiletargeting:*region*:*accountId*:templates/*template-name*/EMAIL

mobiletargeting>DeleteEmailTemplate

Delete a message template for messages that were sent through the email channel.

- URI – [/templates/*template-name*/email](#)
- Method – DELETE
- Resource ARN –
arn:aws:mobiletargeting:*region*:*accountId*:templates/*template-name*/EMAIL

mobiletargeting:GetPushTemplate

Retrieve information about a message template for messages that are sent through a push notification channel.

- URI – [/templates/*template-name*/push](#)
- Method – GET
- Resource ARN –
arn:aws:mobiletargeting:*region*:*accountId*:templates/*template-name*/PUSH

mobiletargeting>CreatePushTemplate

Create a message template for messages that are sent through a push notification channel.

- URI – [/templates/*template-name*/push](#)

- Method – POST
- Resource ARN –
`arn:aws:mobiletargeting:region:accountId:templates/template-name/PUSH`

mobiletargeting:UpdatePushTemplate

Update an existing message template for messages that are sent through a push notification channel.

- URI – [/templates/template-name/push](#)
- Method – PUT
- Resource ARN –
`arn:aws:mobiletargeting:region:accountId:templates/template-name/PUSH`

mobiletargeting>DeletePushTemplate

Delete a message template for messages that were sent through a push notification channel.

- URI – [/templates/template-name/push](#)
- Method – DELETE
- Resource ARN –
`arn:aws:mobiletargeting:region:accountId:templates/template-name/PUSH`

mobiletargeting:GetSmsTemplate

Retrieve information about a message template for messages that are sent through the SMS channel.

- URI – [/templates/template-name/sms](#)
- Method – GET
- Resource ARN –
`arn:aws:mobiletargeting:region:accountId:templates/template-name/SMS`

mobiletargeting>CreateSmsTemplate

Create a message template for messages that are sent through the SMS channel.

- URI – [/templates/template-name/sms](#)

- Method – POST
- Resource ARN –
`arn:aws:mobiletargeting:region:accountId:templates/template-name/SMS`

mobiletargeting:UpdateSmsTemplate

Update an existing message template for messages that are sent through the SMS channel.

- URI – [/templates/*template-name*/sms](#)
- Method – PUT
- Resource ARN –
`arn:aws:mobiletargeting:region:accountId:templates/template-name/SMS`

mobiletargeting>DeleteSmsTemplate

Delete a message template for messages that were sent through the SMS channel.

- URI – [/templates/*template-name*/sms](#)
- Method – DELETE
- Resource ARN –
`arn:aws:mobiletargeting:region:accountId:templates/template-name/SMS`

mobiletargeting:GetVoiceTemplate

Retrieve information about a message template for messages that are sent through the voice channel.

- URI – [/templates/*template-name*/voice](#)
- Method – GET
- Resource ARN –
`arn:aws:mobiletargeting:region:accountId:templates/template-name/VOICE`

mobiletargeting>CreateVoiceTemplate

Create a message template for messages that are sent through the voice channel.

- URI – [/templates/*template-name*/voice](#)
- Method – POST
- Resource ARN –
`arn:aws:mobiletargeting:region:accountId:templates/template-name/VOICE`

mobiletargeting:UpdateVoiceTemplate

Update an existing message template for messages that are sent through the voice channel.

- URI – [/templates/*template-name*/voice](#)
- Method – PUT
- Resource ARN –
arn:aws:mobiletargeting:*region*:*accountId*:templates/*template-name*/VOICE

mobiletargeting>DeleteVoiceTemplate

Delete a message template for messages that were sent through the voice channel.

- URI – [/templates/*template-name*/voice](#)
- Method – DELETE
- Resource ARN –
arn:aws:mobiletargeting:*region*:*accountId*:templates/*template-name*/VOICE

Messages

The following permissions are related to sending messages and push notifications from your Amazon Pinpoint account. You can use the SendMessages and SendUsersMessages operations to send messages to specific endpoints without creating segments and campaigns first.

mobiletargeting:SendMessages

Send a message or push notification to specific endpoints.

- URI – [/apps/*projectId*/messages](#)
- Method – POST
- Resource ARN – arn:aws:mobiletargeting:*region*:*accountId*:apps/*projectId*/messages

mobiletargeting:SendUsersMessages

Send a message or push notification to all the endpoints that are associated with a specific user ID.

- URI – [/apps/*projectId*/users-messages](#)
- Method – POST

- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/messages`

One-time passwords

The following permissions are related to sending and validating one-time passwords (OTPs) in Amazon Pinpoint.

mobiletargeting:SendOTPMessage

Send a text message that contains a one-time password.

- URI – [/apps/projectId/otp](#)
- Method – POST
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/otp`

mobiletargeting:VerifyOTPMessage

Check the validity of a one-time password (OTP) that was generated using the `SendOTPMessage` operation.

- URI – [/apps/projectId/verify-otp](#)
- Method – POST
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/verify-otp`

Phone number validation

The following permissions are related to using the phone number validation service in Amazon Pinpoint.

mobiletargeting:PhoneNumberValidate

Retrieve information about a phone number.

- URI – [/phone/number/validate](#)
- Method – POST
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:phone/number/validate`

Projects

The following permissions are related to managing projects in your Amazon Pinpoint account. Originally, projects were referred to as *applications*. For the purposes of these operations, an Amazon Pinpoint application is the same as an Amazon Pinpoint project.

mobiletargeting:CreateApp

Create an Amazon Pinpoint project.

- URI – [/apps](#)
- Method – POST
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps`

mobiletargeting>DeleteApp

Delete an Amazon Pinpoint project.

- URI – [/apps/projectId](#)
- Method – DELETE
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId`

mobiletargeting:GetApp

Retrieve information about an Amazon Pinpoint project.

- URI – [/apps/projectId](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId`

mobiletargeting:GetApps

Retrieve information about all the projects that are associated with your Amazon Pinpoint account.

- URI – [/apps](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps`

mobiletargeting:GetApplicationSettings

Retrieve the default settings for an Amazon Pinpoint project.

- URI – [/apps/projectId/settings](#)

- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId`

mobiletargeting:UpdateApplicationSettings

Update the default settings for an Amazon Pinpoint project.

- URI – [/apps/projectId/settings](#)
- Method – PUT
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId`

Recommender models

The following permissions are related to managing Amazon Pinpoint configurations for retrieving and processing recommendation data from recommender models. A *recommender model* is a type of machine learning model that predicts and generates personalized recommendations by finding patterns in data.

mobiletargeting:CreateRecommenderConfiguration

Create an Amazon Pinpoint configuration for a recommender model.

- URI – [/recommenders](#)
- Method – POST
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:recommenders`

mobiletargeting:GetRecommenderConfigurations

Retrieve information about all the recommender model configurations that are associated with your Amazon Pinpoint account.

- URI – [/recommenders](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:recommenders`

mobiletargeting:GetRecommenderConfiguration

Retrieve information about an individual Amazon Pinpoint configuration for a recommender model.

- URI – [/recommenders/recommenderId](#)
- Method – GET

- Resource ARN –
arn:aws:mobiletargeting:*region*:*accountId*:recommenders/*recommenderId*

mobiletargeting:UpdateRecommenderConfiguration

Update an Amazon Pinpoint configuration for a recommender model.

- URI – [/recommenders/*recommenderId*](#)
- Method – PUT
- Resource ARN –
arn:aws:mobiletargeting:*region*:*accountId*:recommenders/*recommenderId*

mobiletargeting>DeleteRecommenderConfiguration

Delete an Amazon Pinpoint configuration for a recommender model.

- URI – [/recommenders/*recommenderId*](#)
- Method – DELETE
- Resource ARN –
arn:aws:mobiletargeting:*region*:*accountId*:recommenders/*recommenderId*

Segments

The following permissions are related to managing segments in your Amazon Pinpoint account. In Amazon Pinpoint, *segments* are groups of recipients for your campaigns that share certain attributes that you define.

mobiletargeting>CreateSegment

Create a segment. To allow a user to create a segment by importing endpoint data from outside Amazon Pinpoint, allow the `mobiletargeting:CreateImportJob` action.

- URI – [/apps/*projectId*/segments](#)
- Method – POST
- Resource ARN – arn:aws:mobiletargeting:*region*:*accountId*:apps/*projectId*

mobiletargeting>DeleteSegment

Delete a segment.

- URI – [/apps/*projectId*/segments/*segmentId*](#)
- Method – DELETE

- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/segments/segmentId`

mobiletargeting:GetSegment

Retrieve information about a specific segment.

- URI – [/apps/projectId/segments/segmentId](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/segments/segmentId`

mobiletargeting:GetSegmentExportJobs

Retrieve information about jobs that export endpoint definitions for a segment.

- URI – [/apps/projectId/segments/segmentId/jobs/export](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/segments/segmentId/jobs/export`

mobiletargeting:GetSegments

Retrieve information about all the segments for a project.

- URI – [/apps/projectId/segments](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId`

mobiletargeting:GetSegmentImportJobs

Retrieve information about jobs that create segments by importing endpoint definitions from Amazon S3.

- URI – [/apps/projectId/segments/segmentId/jobs/import](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/segments/segmentId`

mobiletargeting:GetSegmentVersion

Retrieve information about a specific segment version.

- URI – [/apps/projectId/segments/segmentId/versions/versionId](#)
- Method – GET

- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/segments/segmentId`

mobiletargeting:GetSegmentVersions

Retrieve information about the current and prior versions of a segment.

- URI – [/apps/projectId/segments/segmentId/versions](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/segments/segmentId`

mobiletargeting:UpdateSegment

Update a specific segment.

- URI – [/apps/projectId/segments/segmentId](#)
- Method – PUT
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/segments/segmentId`

Tags

The following permissions are related to viewing and managing tags for Amazon Pinpoint resources.

mobiletargeting:ListTagsForResource

Retrieve information about the tags that are associated with a project, campaign, message template, or segment.

- URI – [/tags/resource-arn](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:*`

mobiletargeting:TagResource

Add one or more tags to a project, campaign, message template, or segment.

- URI – [/tags/resource-arn](#)
- Method – POST

- Resource ARN – `arn:aws:mobiletargeting:region:accountId:*`

mobiletargeting:UntagResource

Remove one or more tags from a project, campaign, message template, or segment.

- URI – [/tags/resource-arn](#)
- Method – DELETE
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:*`

Users

The following permissions are related to managing users. In Amazon Pinpoint, *users* correspond to individuals who receive messages from you. A single user might be associated with more than one endpoint.

mobiletargeting>DeleteUserEndpoints

Delete all the endpoints that are associated with a user ID.

- URI – [/apps/projectId/users/userId](#)
- Method – DELETE
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/users/userId`

mobiletargeting:GetUserEndpoints

Retrieve information about all the endpoints that are associated with a user ID.

- URI – [/apps/projectId/users/userId](#)
- Method – GET
- Resource ARN – `arn:aws:mobiletargeting:region:accountId:apps/projectId/users/userId`

Amazon Pinpoint SMS and voice version 1 API actions

This section identifies actions for features that are available from the Amazon Pinpoint SMS and Voice API. This is a supplemental API that provides advanced options for using and managing

the SMS and voice channels in Amazon Pinpoint. To learn more about this API, see the [Amazon Pinpoint SMS and voice API reference](#).

sms-voice:CreateConfigurationSet

Create a configuration set for sending voice messages.

- URI – /sms-voice/configuration-sets
- Method – POST
- Resource ARN – Not available. Use *.

sms-voice>DeleteConfigurationSet

Delete a configuration set for sending voice messages.

- URI – /sms-voice/configuration-sets/*ConfigurationSetName*
- Method – DELETE
- Resource ARN – Not available. Use *.

sms-voice:GetConfigurationSetEventDestinations

Retrieve information about a configuration set and the event destinations that it contains.

- URI – /sms-voice/configuration-sets/*ConfigurationSetName*/event-destinations
- Method – GET
- Resource ARN – Not available. Use *.

sms-voice:CreateConfigurationSetEventDestination

Create an event destination for voice events.

- URI – /sms-voice/configuration-sets/*ConfigurationSetName*/event-destinations
- Method – POST
- Resource ARN – Not available. Use *.

sms-voice:UpdateConfigurationSetEventDestination

Update an event destination for voice events.

- URI – /sms-voice/configuration-sets/*ConfigurationSetName*/event-destinations/*EventDestinationName*
- Method – PUT
- Resource ARN – Not available. Use *.

sms-voice:DeleteConfigurationSetEventDestination

Delete an event destination for voice events.

- URI – /sms-voice/configuration-sets/*ConfigurationSetName*/event-destinations/*EventDestinationName*
- Method – DELETE
- Resource ARN – Not available. Use *.

sms-voice:SendVoiceMessage

Create and send voice messages.

- URI – /sms-voice/voice/message
- Method – POST
- Resource ARN – Not available. Use *.

AWS End User Messaging SMS and voice version 2 API actions

This section identifies actions for features that are available from the Amazon Pinpoint SMS and Voice API. For the Amazon Pinpoint SMS and Voice API, there is a supplemental API that provides advanced options for using and managing the SMS and voice channels. For a complete list of actions available in version 2, see the [AWS End User Messaging SMS and Voice API version 2 API Reference](#).

sms-voice:AssociateOriginationIdentity

Associate the specified origination identity with a pool.

- Resource ARN – arn:aws:sms-voice:*region*:*accountId*:pool/*poolId*
- Resource ARN – arn:aws:sms-voice:*region*:*accountId*:phone-number/*phoneNumberId*

- Resource ARN – `arn:aws:sms-voice:region:accountId:sender-id/senderId/isoCountyCode`

sms-voice:CreateConfigurationSet

Create a new configuration set.

- Resource ARN – `arn:aws:sms-voice:region:accountId:configuration-set/configurationSetName`

sms-voice:CreateEventDestination

Create a new event destination in a configuration set.

- Resource ARN – `arn:aws:sms-voice:region:accountId:configuration-set/configurationSetName`

sms-voice:CreateOptOutList

Create a new opt-out list.

- Resource ARN – `arn:aws:sms-voice:region:accountId:opt-out-list/optOutListName`

sms-voice:CreatePool

Create a new pool and associates the specified origination identity to the pool.

- Resource ARN – `arn:aws:sms-voice:region:accountId:pool/poolId`
- Resource ARN – `arn:aws:sms-voice:region:accountId:phone-number/phoneNumberId`
- Resource ARN – `arn:aws:sms-voice:region:accountId:sender-id/senderId/isoCountyCode`

sms-voice>DeleteConfigurationSet

Delete an existing configuration set.

- Resource ARN – `arn:aws:sms-voice:region:accountId:configuration-set/configurationSetName`

sms-voice>DeleteDefaultMessageType

Delete an existing default message type on a configuration set.

- Resource ARN – `arn:aws:sms-voice:region:accountId:configuration-set/configurationSetName`

sms-voice:DeleteDefaultSenderId

Delete an existing default Sender ID on a configuration set.

- Resource ARN – `arn:aws:sms-voice:region:accountId:sender-id/configuration-set/configurationSetName`

sms-voice:DeleteEventDestination

Delete an existing event destination.

- Resource ARN – `arn:aws:sms-voice:region:accountId:configuration-set/configurationSetName`

sms-voice:DeleteKeyword

Delete an existing keyword from an origination phone number or pool.

- Resource ARN – `arn:aws:sms-voice:region:accountId:pool/poolId`
- Resource ARN – `arn:aws:sms-voice:region:accountId:phone-number/phoneNumberId`

sms-voice:DeleteOptedOutNumber

Delete an existing opted out destination phone number from the specified opt-out list.

- Resource ARN – `arn:aws:sms-voice:region:accountId:opt-out-list/optOutListName`

sms-voice:DeleteOptOutList

Delete an existing opt-out list. All opted out phone numbers in the opt-out list are deleted.

- Resource ARN – `arn:aws:sms-voice:region:accountId:opt-out-list/optOutListName`

sms-voice:DeletePool

Delete an existing pool.

- Resource ARN – `arn:aws:sms-voice:region:accountId:pool/poolId`

sms-voice:DeleteTextMessageSpendLimitOverride

Delete an account-level monthly spending limit override for sending text messages.

- Resource ARN – Not available. Use *.

sms-voice:DeleteVoiceMessageSpendLimitOverride

Delete an account-level monthly spend limit override for sending voice messages.

- Resource ARN – Not available. Use *.

sms-voice:DescribeAccountAttributes

Describe attributes of your AWS account.

- Resource ARN – Not available. Use *.

sms-voice:DescribeAccountLimits

Describe the current resource quotas for your account.

- Resource ARN – Not available. Use *.

sms-voice:DescribeConfigurationSets

Describe the specified configuration sets or all in your account.

- Resource ARN – `arn:aws:sms-voice:region:accountId:configuration-set/configurationSetName`

sms-voice:DescribeKeywords

Describe the specified keywords or all keywords on your origination phone number or pool.

- Resource ARN – `arn:aws:sms-voice:region:accountId:pool/poolId`
- Resource ARN – `arn:aws:sms-voice:region:accountId:phone-number/phoneNumberId`

sms-voice:DescribeOptedOutNumbers

Describe the specified opted out destination numbers or all opted out destination numbers in an opt-out list.

- Resource ARN – `arn:aws:sms-voice:region:accountId:opt-out-list/optOutListName`

sms-voice:DescribeOptOutLists

Describe the specified opt-out list or all opt-out lists in your account.

- Resource ARN – `arn:aws:sms-voice:region:accountId:opt-out-list/optOutListName`

sms-voice:DescribePhoneNumbers

Describe the specified origination phone number, or all the phone numbers in your account.

- Resource ARN – `arn:aws:sms-voice:region:accountId:phone-number/phoneNumberId`

sms-voice:DescribePools

Retrieve the specified pools or all pools associated with your AWS account.

- Resource ARN – `arn:aws:sms-voice:region:accountId:pool/poolId`

sms-voice:DescribeSenderId

Describe the specified Sender IDs or all Sender IDs associated with your AWS account.

- Resource ARN – `arn:aws:sms-voice:region:accountId:sender-id/senderId/isoCountryCode`

sms-voice:DescribeSpendLimits

Describe the current Amazon Pinpoint monthly spend limits for sending voice and text messages.

- Resource ARN – Not available. Use *.

sms-voice:DisassociateOriginationIdentity

Remove the specified origination identity from an existing pool.

- Resource ARN – `arn:aws:sms-voice:region:accountId:pool/poolId`
- Resource ARN – `arn:aws:sms-voice:region:accountId:phone-number/phoneNumberId`
- Resource ARN – `arn:aws:sms-voice:region:accountId:sender-id/senderId/isoCountryCode`

sms-voice:ListPoolOriginationIdentities

Show the origination phone numbers in a pool.

- Resource ARN – `arn:aws:sms-voice:region:accountId:pool/poolId`

sms-voice:ListTagsForResource

List the tags associated with a resource.

- Resource ARN – `arn:aws:sms-voice:region:accountId:configuration-set/configurationSetName`
- Resource ARN – `arn:aws:sms-voice:region:accountId:opt-out-list/optOutListName`

- Resource ARN – `arn:aws:sms-voice:region:accountId:phone-number/phoneNumberId`
- Resource ARN – `arn:aws:sms-voice:region:accountId:pool/poolId`
- Resource ARN – `arn:aws:sms-voice:region:accountId:sender-id/senderId/isoCountryCode`

sms-voice:PutKeyword

Add or update a keyword on an origination phone number or pool.

- Resource ARN – `arn:aws:sms-voice:region:accountId:phone-number/phoneNumberId`
- Resource ARN – `arn:aws:sms-voice:region:accountId:pool/poolId`

sms-voice:PutOptedOutNumber

Add a destination phone number to an opt-out list.

- Resource ARN – `arn:aws:sms-voice:region:accountId:opt-out-list/optOutListName`

sms-voice:ReleasePhoneNumber

Remove an origination phone number from your Amazon Pinpoint account.

- Resource ARN – `arn:aws:sms-voice:region:accountId:phone-number/phoneNumberId`

sms-voice:RequestPhoneNumber

Request to add an origination phone number to your account.

- Resource ARN – `arn:aws:sms-voice:region:accountId:phone-number/phoneNumberId`
- Resource ARN – `arn:aws:sms-voice:region:accountId:pool/poolId`

sms-voice:SendTextMessage

Send an SMS message.

- Resource ARN – `arn:aws:sms-voice:region:accountId:phone-number/phoneNumberId`
- Resource ARN – `arn:aws:sms-voice:region:accountId:pool/poolId`
- Resource ARN – `arn:aws:sms-voice:region:accountId:sender-id/senderId/isoCountryCode`

sms-voice:SendVoiceMessage

Send a voice message.

- Resource ARN – arn:aws:sms-voice:*region*:*accountId*:phone-number/*phoneNumberId*
- Resource ARN – arn:aws:sms-voice:*region*:*accountId*:pool/*poolId*

sms-voice:SetDefaultMessageType

Set the default message type for SMS messages.

- Resource ARN – arn:aws:sms-voice:*region*:*accountId*:configuration-set/*configurationSetName*

sms-voice:SetDefaultSenderId

Set the default Sender ID value for voice messages.

- Resource ARN – arn:aws:sms-voice:*region*:*accountId*:configuration-set/*configurationSetName*

sms-voice:SetTextMessageSpendLimitOverride

Set a monthly spending limit for SMS messages.

- Resource ARN – Not available. Use *.

sms-voice:SetVoiceMessageSpendLimitOverride

Set a monthly spending limit for voice messages.

- Resource ARN – Not available. Use *.

sms-voice:TagResource

Add a tag to a resource.

- Resource ARN – arn:aws:sms-voice:*region*:*accountId*:configuration-set/*configurationSetName*
- Resource ARN – arn:aws:sms-voice:*region*:*accountId*:opt-out-list/*optOutListName*
- Resource ARN – arn:aws:sms-voice:*region*:*accountId*:phone-number/*phoneNumberId*
- Resource ARN – arn:aws:sms-voice:*region*:*accountId*:pool/*poolId*
- Resource ARN – arn:aws:sms-voice:*region*:*accountId*:sender-id/*senderId/isoCountryCode*

sms-voice:UntagResource

Remove tags from a resource.

- Resource ARN – `arn:aws:sms-voice:region:accountId:configuration-set/configurationSetName`
- Resource ARN – `arn:aws:sms-voice:region:accountId:opt-out-list/optOutListName`
- Resource ARN – `arn:aws:sms-voice:region:accountId:phone-number/phoneNumberId`
- Resource ARN – `arn:aws:sms-voice:region:accountId:pool/poolId`
- Resource ARN – `arn:aws:sms-voice:region:accountId:sender-id/senderId/isoCountryCode`

sms-voice:UpdateEventDestination

Update an existing event destination.

- Resource ARN – `arn:aws:sms-voice:region:accountId:configuration-set/configurationSetName`

sms-voice:UpdatePhoneNumber

Update the configuration of an origination phone number.

- Resource ARN – `arn:aws:sms-voice:region:accountId:phone-number/phoneNumberId`

sms-voice:UpdatePool

Update an existing phone number pool.

- Resource ARN – `arn:aws:sms-voice:region:accountId:pool/poolId`

Amazon Pinpoint identity-based policy examples

By default, IAM users and roles don't have permission to create or modify Amazon Pinpoint resources, and they also can't perform tasks using the AWS Management Console, AWS CLI, or an AWS API. An IAM administrator must create IAM policies that grant users and roles permission to perform specific API operations on the resources that they need. The administrator must then attach those policies to the IAM users or groups that require those permissions.

To learn how to create an IAM identity-based policy using these example JSON policy documents, see [Creating IAM policies](#) in the *IAM User Guide*.

Topics

- [Policy best practices](#)
- [Using the Amazon Pinpoint console](#)
- [Example: Accessing a single Amazon Pinpoint project](#)
- [Example: Viewing Amazon Pinpoint resources based on tags](#)
- [Example: Allowing users to view their own permissions](#)
- [Examples: Providing access to Amazon Pinpoint API actions](#)
- [Examples: Providing access to Amazon Pinpoint SMS and voice API actions](#)
- [Example: Restricting Amazon Pinpoint project access to specific IP addresses](#)
- [Example: Restricting Amazon Pinpoint access based on tags](#)
- [Example: Allowing Amazon Pinpoint to send email using identities that were verified in Amazon SES](#)

Policy best practices

Identity-based policies determine whether someone can create, access, or delete Amazon Pinpoint resources in your account. These actions can incur costs for your AWS account. When you create or edit identity-based policies, follow these guidelines and recommendations:

- **Get started with AWS managed policies and move toward least-privilege permissions** – To get started granting permissions to your users and workloads, use the *AWS managed policies* that grant permissions for many common use cases. They are available in your AWS account. We recommend that you reduce permissions further by defining AWS customer managed policies that are specific to your use cases. For more information, see [AWS managed policies](#) or [AWS managed policies for job functions](#) in the *IAM User Guide*.
- **Apply least-privilege permissions** – When you set permissions with IAM policies, grant only the permissions required to perform a task. You do this by defining the actions that can be taken on specific resources under specific conditions, also known as *least-privilege permissions*. For more information about using IAM to apply permissions, see [Policies and permissions in IAM](#) in the *IAM User Guide*.
- **Use conditions in IAM policies to further restrict access** – You can add a condition to your policies to limit access to actions and resources. For example, you can write a policy condition to


```

    {
      "Sid": "UseConsole",
      "Effect": "Allow",
      "Action": [
        "mobiletargeting:Get*",
        "mobiletargeting:List*"
      ],
      "Resource": "arn:aws:mobiletargeting:us-east-1:111122223333:*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "firehose:ListDeliveryStreams",
        "iam:ListRoles",
        "kinesis:ListStreams",
        "s3:List*",
        "ses:Describe*",
        "ses:Get*",
        "ses:List*",
        "sns:ListTopics"
      ],
      "Resource": "*"
    }
  ]
}

```

In the preceding policy example, replace *region* with the name of an AWS Region, and replace *accountId* with your AWS account ID.

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, only allow access to the actions that match the API operation that they're trying to perform.

Example: Accessing a single Amazon Pinpoint project

You can also create read-only policies that only provide access to specific projects. The following example policy lets users sign in to the console and view a list of projects. It also lets users view information about related resources for other AWS services that the Amazon Pinpoint console depends on, such as Amazon SES, IAM, and Amazon Kinesis. However, the policy lets users only view additional information about the project that's specified in the policy. You can modify this policy to allow access to additional projects or AWS Regions.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ViewProject",
      "Effect": "Allow",
      "Action": "mobiletargeting:GetApps",
      "Resource": "arn:aws:mobiletargeting:us-east-1:111122223333:*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "mobiletargeting:Get*",
        "mobiletargeting:List*"
      ],
      "Resource": [
        "arn:aws:mobiletargeting:us-east-1:111122223333:apps/projectId",
        "arn:aws:mobiletargeting:us-east-1:111122223333:apps/projectId/"
      ],
      "Resource": "arn:aws:mobiletargeting:us-east-1:111122223333:reports"
    }
  ],
  {
    "Effect": "Allow",
    "Action": [
      "ses:Get*",
      "kinesis:ListStreams",
      "firehose:ListDeliveryStreams",
      "iam:ListRoles",
      "ses:List*",
      "sns:ListTopics",
      "ses:Describe*",
      "s3:List*"
    ],
    "Resource": "*"
  }
]
}

```

In the preceding example, replace *region* with the name of an AWS Region, replace *accountId* with your AWS account ID, and replace *projectId* with the ID of the Amazon Pinpoint project that you want to provide access to.

Similarly, you can create policies that grant an IAM user in your AWS account with limited write access to a specific Amazon Pinpoint project. In this case, you want to allow the user to view, add, and update project components, such as segments and campaigns, but not delete any components.

In addition to granting permissions for `mobiletargeting:Get` and `mobiletargeting:List` actions, create a policy that grants permissions for the following actions: `mobiletargeting:Create`; `mobiletargeting:Update`; and `mobiletargeting:Put`. These are the additional permissions required to create and manage most project components. For example:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "LimitedWriteProject",
      "Effect": "Allow",
      "Action": "mobiletargeting:GetApps",
      "Resource": "arn:aws:mobiletargeting:us-east-1:111122223333:*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "mobiletargeting:Get*",
        "mobiletargeting:List*",
        "mobiletargeting:Create*",
        "mobiletargeting:Update*",
        "mobiletargeting:Put*"
      ],
      "Resource": [
        "arn:aws:mobiletargeting:us-east-1:111122223333:apps/810c7aab86d42fb2b56c8c966example",
        "arn:aws:mobiletargeting:us-east-1:111122223333:apps/810c7aab86d42fb2b56c8c966example/*",
        "arn:aws:mobiletargeting:us-east-1:111122223333:reports"
      ]
    }
  ]
}
```

```

    },
    {
      "Effect": "Allow",
      "Action": [
        "ses:Get*",
        "kinesis:ListStreams",
        "firehose:ListDeliveryStreams",
        "iam:ListRoles",
        "ses:List*",
        "sns:ListTopics",
        "ses:Describe*",
        "s3:List*"
      ],
      "Resource": "*"
    }
  ]
}

```

Example: Viewing Amazon Pinpoint resources based on tags

You can use conditions in an identity-based policy to control access to Amazon Pinpoint resources based on tags. This example policy shows how you might create this kind of policy to allow viewing Amazon Pinpoint resources. However, permission is granted only if the `Owner` resource tag has the value of that user's user name. This policy also grants the permissions necessary to complete this action on the console.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ListResources",
      "Effect": "Allow",
      "Action": [
        "mobiletargeting:Get*",
        "mobiletargeting:List*"
      ],
      "Resource": "*"
    },
    {

```

```

        "Sid": "ViewResourceIfOwner",
        "Effect": "Allow",
        "Action": [
            "mobiletargeting:Get*",
            "mobiletargeting:List*"
        ],
        "Resource": "arn:aws:mobiletargeting:*:*:*",
        "Condition": {
            "StringEquals": {"aws:ResourceTag/Owner": "${aws:username}"}
        }
    }
}

```

You can attach this type of policy to the IAM users in your account. If a user named `richard-roe` attempts to view an Amazon Pinpoint resource, the resource must be tagged `Owner=richard-roe` or `owner=richard-roe`. Otherwise, he is denied access. The condition tag key `Owner` matches both `Owner` and `owner` because condition key names are not case-sensitive. For more information, see [IAM JSON policy elements: Condition](#) in the *IAM User Guide*.

Example: Allowing 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:ListGroupsWithUser",
        "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": "*"
    }
  ]
}

```

Examples: Providing access to Amazon Pinpoint API actions

This section provides example policies that allow you to access features that are available from the Amazon Pinpoint API, which is the primary API for Amazon Pinpoint. To learn more about this API, see the [Amazon Pinpoint API Reference](#).

Read-only access

The following example policy allows you read-only access to all the resources in your Amazon Pinpoint account in a specific AWS Region.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ViewAllResources",
      "Effect": "Allow",
      "Action": [
        "mobiletargeting:Get*",
        "mobiletargeting:List*"
      ],
      "Resource": "arn:aws:mobiletargeting:us-east-1:111122223333:*"
    }
  ]
}

```

```

    }
  ]
}

```

In the preceding example, replace *region* with the name of an AWS Region, and replace *accountId* with your AWS account ID.

Administrator access

The following example policy allows you full access to all Amazon Pinpoint actions and resources in your Amazon Pinpoint account in all AWS Regions:

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "FullAccess",
      "Effect": "Allow",
      "Action": [
        "mobiletargeting:*"
      ],
      "Resource": "arn:aws:mobiletargeting:*:111122223333:*"
    }
  ]
}

```

In the preceding example, replace *accountId* with your AWS account ID.

Examples: Providing access to Amazon Pinpoint SMS and voice API actions

This section provides example policies that allow you to access features that are available from the Amazon Pinpoint SMS and Voice API. This is a supplemental API that provides advanced options for using and managing the SMS and voice channels in Amazon Pinpoint. To learn more about this API, see the [Amazon Pinpoint SMS and voice API reference](#).

Read-only access

The following example policy allows you read-only access to all Amazon Pinpoint SMS and Voice API actions and resources in your AWS account in all AWS Regions:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ViewAllResources",
      "Effect": "Allow",
      "Action": [
        "sms-voice:Get*",
        "sms-voice:List*"
      ],
      "Resource": "*"
    }
  ]
}
```

Administrator access

The following example policy allows you full access to all Amazon Pinpoint SMS and Voice API actions and resources in your AWS account in all AWS Regions:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "FullAccess",
      "Effect": "Allow",
      "Action": [
        "sms-voice:*"
      ],
      "Resource": "*"
    }
  ]
}
```

```

        "Resource": "*"
    }
]
}

```

Example: Restricting Amazon Pinpoint project access to specific IP addresses

The following example policy grants permissions to any user to perform any Amazon Pinpoint action on a specified project (*projectId*). However, the request must originate from the range of IP addresses that are specified in the condition.

The condition in this statement identifies the `54.240.143.*` range of allowed Internet Protocol version 4 (IPv4) addresses, with one exception: `54.240.143.188`. The Condition block uses the `IpAddress` and `NotIpAddress` conditions and the `aws:SourceIp` condition key, which is an AWS-wide condition key. For more information about these condition keys, see [Specifying conditions in a policy](#) *IAM User Guide*. The `aws:SourceIp` IPv4 values use standard CIDR notation. For more information, see [IP address condition operators](#) in the *IAM User Guide*.

JSON

```

{
  "Version": "2012-10-17",
  "Id": "AMZPinpointPolicyId1",
  "Statement": [
    {
      "Sid": "IPAllow",
      "Effect": "Allow",
      "Principal": "*",
      "Action": "mobiletargeting:*",
      "Resource": [
        "arn:aws:mobiletargeting:*:*:apps/projectId",
        "arn:aws:mobiletargeting:*:*:apps/projectId/*"
      ],
      "Condition": {
        "IpAddress": {"aws:SourceIp": "54.240.143.0/24"},
        "NotIpAddress": {"aws:SourceIp": "54.240.143.188/32"}
      }
    }
  ]
}

```

Example: Restricting Amazon Pinpoint access based on tags

The following example policy grants you permissions to perform any Amazon Pinpoint action on a specified project (*projectId*). However, permissions are granted only if the request originates from a user whose name is a value in the `Owner` resource tag for the project, as specified in the condition.

The `Condition` block uses the `StringEquals` condition and the `aws:ResourceTag/${TagKey}` condition key. For more information about conditions and condition keys, see [Bucket policy examples using condition keys](#) in the *IAM User Guide*.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ModifyResourceIfOwner",
      "Effect": "Allow",
      "Action": "mobiletargeting:*",
      "Resource": [
        "arn:aws:mobiletargeting:*:*:apps/projectId",
        "arn:aws:mobiletargeting:*:*:apps/projectId/*"
      ],
      "Condition": {
        "StringEquals": {"aws:ResourceTag/Owner": "${aws:username}"}
      }
    }
  ]
}
```

Example: Allowing Amazon Pinpoint to send email using identities that were verified in Amazon SES

When you verify an email identity (such as an email address or domain) through the Amazon Pinpoint console, that identity is automatically configured so that it can be used by both Amazon Pinpoint and Amazon SES. However, if you verify an email identity through Amazon SES, and you want to use that identity with Amazon Pinpoint, you must apply a policy to that identity.

The following example policy grants Amazon Pinpoint permission to send email using an email identity that was verified through Amazon SES.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "PinpointEmail",
      "Effect": "Allow",
      "Principal": {
        "Service": "pinpoint.amazonaws.com"
      },
      "Action": "ses:*",
      "Resource": "arn:aws:ses:us-east-1:111122223333:identity/emailId",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "111122223333"
        },
        "ArnLike": {
          "aws:SourceArn": "arn:aws:mobiletargeting:us-east-1:111122223333:apps/*"
        }
      }
    }
  ]
}
```

If you use Amazon Pinpoint in the AWS GovCloud (US-West) Region, use the following policy example instead:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "PinpointEmail",
      "Effect": "Allow",
```

```
    "Principal": {
      "Service": "pinpoint.amazonaws.com"
    },
    "Action": "ses:*",
    "Resource": "arn:aws-us-gov:ses:us-gov-
west-1:111122223333:identity/emailId",
    "Condition": {
      "StringEquals": {
        "aws:SourceAccount": "111122223333"
      },
      "ArnLike": {
        "aws:SourceArn": "arn:aws-us-gov:mobiletargeting:us-gov-
west-1:111122223333:apps/*"
      }
    }
  }
]
```

IAM roles for common Amazon Pinpoint tasks

An [IAM role](#) is an AWS Identity and Access Management (IAM) identity that you can create in your AWS account and grant specific permissions. An IAM role is similar to an IAM user, in that it's an AWS identity with permissions policies that determine what the identity can and can't do in AWS. However, instead of being uniquely associated with one person, a role is intended to be assumable by anyone who needs it.

Also, a role doesn't have standard long-term credentials such as a password or access keys associated with it. Instead, it provides temporary security credentials for a session. You can use IAM roles to delegate access to users, apps, applications, or services that don't normally have access to your AWS resources.

For these reasons, you can use IAM roles to integrate Amazon Pinpoint with certain AWS services and resources for your account. For example, you might want to allow Amazon Pinpoint to access endpoint definitions that you store in an Amazon Simple Storage Service (Amazon S3) bucket and want to use for segments. Or you might want to allow Amazon Pinpoint to stream event data to an Amazon Kinesis stream for your account. Similarly, you might want to use IAM roles to allow web or mobile apps to register endpoints or report usage data for Amazon Pinpoint projects, without embedding AWS keys in the apps (where they can be difficult to rotate and users can potentially extract them).

For these scenarios, you can delegate access to Amazon Pinpoint by using IAM roles. This section explains and provides examples of common Amazon Pinpoint tasks that use IAM roles to work with other AWS services. For information about using IAM roles with web and mobile apps, see [Providing access to externally authenticated users \(identity federation\)](#) in the *IAM User Guide*.

Topics

- [IAM role for importing endpoints or segments](#)
- [IAM role for exporting endpoints or segments](#)
- [Retrieving recommendations from Amazon Personalize](#)
- [IAM role for streaming events to Kinesis](#)
- [IAM role for streaming email events to Firehose](#)

IAM role for importing endpoints or segments

With Amazon Pinpoint, you can define a user segment by importing endpoint definitions from an Amazon Simple Storage Service (Amazon S3) bucket in your AWS account. Before you import, you must delegate the required permissions to Amazon Pinpoint. To do this, you create an AWS Identity and Access Management (IAM) role and attach the following policies to the role:

- The `AmazonS3ReadOnlyAccess` AWS managed policy. This policy is created and managed by AWS, and it grants read-only access to your Amazon S3 bucket.
- A trust policy that allows Amazon Pinpoint to assume the role.

After you create the role, you can use Amazon Pinpoint to import segments from an Amazon S3 bucket. For information about creating the bucket, creating endpoint files, and importing a segment by using the console, see [Importing segments](#) in the *Amazon Pinpoint User Guide*. For an example of how to import a segment programmatically by using the AWS SDK for Java, see [Importing segments](#) in the *Amazon Pinpoint Developer Guide*.

Attaching the trust policy

To allow Amazon Pinpoint to assume the IAM role and perform the actions allowed by the `AmazonS3ReadOnlyAccess` policy, attach the following trust policy to the role:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowUserToImportEndpointDefinitions",
      "Effect": "Allow",
      "Principal": {
        "Service": "pinpoint.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

Creating the IAM role (AWS CLI)

Complete the following steps to create the IAM role by using the AWS Command Line Interface (AWS CLI). If you haven't installed the AWS CLI, see [Install or update to the latest version of AWS CLI](#) in the *AWS Command Line Interface User Guide*.

To create the IAM role by using the AWS CLI

1. Create a JSON file that contains the trust policy for your role, and save the file locally. You can copy the trust policy provided in this topic.
2. At the command line, use the [create-role](#) command to create the role and attach the trust policy:

```
aws iam create-role --role-name PinpointSegmentImport --assume-role-policy-document
file://PinpointImportTrustPolicy.json
```

Following the `file://` prefix, specify the path to the JSON file that contains the trust policy.

After you run this command, you will see an output that's similar to the following in your terminal:

3. Use the [attach-role-policy](#) command to attach the AmazonS3ReadOnlyAccess AWS managed policy to the role:

```
aws iam attach-role-policy --policy-arn arn:aws:iam::aws:policy/  
AmazonS3ReadOnlyAccess --role-name PinpointSegmentImport
```

IAM role for exporting endpoints or segments

You can obtain a list of endpoints by creating an export job. When you create an export job, you must specify a project ID, and you can optionally specify a segment ID. Amazon Pinpoint then exports a list of the endpoints associated with the project or segment to an Amazon Simple Storage Service (Amazon S3) bucket. The resulting file contains a JSON-formatted list of endpoints and their attributes, such as channel, address, opt-in/opt-out status, creation date, and endpoint ID.

To create an export job, you must configure an IAM role that allows Amazon Pinpoint to write to an Amazon S3 bucket. The process of configuring the role consists of two steps:

1. Create an IAM policy that allows an entity (in this case, Amazon Pinpoint) to write to a specific Amazon S3 bucket.
2. Create an IAM role and attach the policy to it.

This topic contains procedures for completing both of these steps. These procedures assume that you've already created an Amazon S3 bucket, and a folder in that bucket, for storing exported segments. For information about creating buckets, see [Create a bucket](#) in the *Amazon Simple Storage Service User Guide*.

These procedures also assume that you've already installed and configured the AWS Command Line Interface (AWS CLI). For information about setting up the AWS CLI, see [Get started with the AWS CLI](#) in the *AWS Command Line Interface User Guide*.

Step 1: Create the IAM policy

An IAM policy defines the permissions for an entity, such as an identity or resource. To create a role for exporting Amazon Pinpoint endpoints, you must create a policy that grants permission to write to a specific folder in a specific Amazon S3 bucket. The following policy example follows the security practice of granting least privilege, meaning it grants only the permissions that are required to perform a single task.

To create the IAM policy

1. In a text editor, create a new file. Paste the following code into the file:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowUserToSeeBucketListInTheConsole",
      "Action": [
        "s3:ListAllMyBuckets",
        "s3:GetBucketLocation"
      ],
      "Effect": "Allow",
      "Resource": [ "arn:aws:s3:::*" ]
    },
    {
      "Sid": "AllowRootAndHomeListingOfBucket",
      "Action": [
        "s3:ListBucket"
      ],
      "Effect": "Allow",
      "Resource": [ "arn:aws:s3:::amzn-s3-demo-bucket-example-bucket" ],
      "Condition": {
        "StringEquals": {
          "s3:delimiter": [ "/" ],
          "s3:prefix": [
            "",
            Exports/
          ]
        }
      }
    },
    {
      "Sid": "AllowListingOfUserFolder",
      "Action": [
        "s3:ListBucket"
      ],
      "Effect": "Allow",
```

```

        "Resource": [ "arn:aws:s3:::amzn-s3-demo-bucket-example-
bucket" ],
        "Condition": {
            "StringLike": {
                "s3:prefix": [
                    "Exports/"
                ]
            }
        },
        {
            "Sid": "AllowAllS3ActionsInUserFolder",
            "Action": [ "s3:*" ],
            "Effect": "Allow",
            "Resource": [ "arn:aws:s3:::amzn-s3-demo-bucket-example-
bucket/Exports/" ]
        }
    ]
}

```

In the preceding code, replace all instances of `amzn-s3-demo-bucket-example-bucket` with the name of the Amazon S3 bucket that contains the folder that you want to export the segment information into. Also, replace all instances of `Exports` with the name of the folder itself.

When you finish, save the file as `s3policy.json`.

- By using the AWS CLI, navigate to the directory where the `s3policy.json` file is located. Then enter the following command to create the policy:

```
aws iam create-policy --policy-name s3ExportPolicy --policy-document
file://s3policy.json
```

If the policy was created successfully, you will see an output similar to the following:

```
{
  "Policy": {
    "CreateDate": "2018-04-11T18:44:34.805Z",
    "IsAttachable": true,
    "DefaultVersionId": "v1",
    "AttachmentCount": 0,
  }
}
```

```
"PolicyId": "ANPAJ2YJQRJCG3EXAMPLE",
"UpdateDate": "2018-04-11T18:44:34.805Z",
"Arn": "arn:aws:iam::123456789012:policy/s3ExportPolicy",
"PolicyName": "s3ExportPolicy",
"Path": "/"
}
}
```

Copy the Amazon Resource Name (ARN) of the policy (arn:aws:iam::123456789012:policy/s3ExportPolicy in the preceding example). In the next section, you must supply this ARN when you create the role.

Note

If you receive a message stating that your account isn't authorized to perform the CreatePolicy operation, then you must attach a policy to your user account that lets you create new IAM policies and roles. For more information, see [Adding and removing IAM identity permissions](#) in the *IAM User Guide*.

Step 2: Create the IAM role

Now that you've created an IAM policy, you can create a role and attach the policy to it. Each IAM role contains a *trust policy*—a set of rules that specifies which entities are allowed to assume the role. In this section, you create a trust policy that allows Amazon Pinpoint to assume the role. Next, you create the role itself, and then attach the policy that you created in the previous section.

To create the IAM role

1. In a text editor, create a new file. Paste the following code into the file:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "pinpoint.amazonaws.com"
      }
    }
  ]
}
```

```
        },
        "Action": "sts:AssumeRole"
    }
]
}
```

Save the file as `trustpolicy.json`.

- By using the AWS CLI, navigate to the directory where the `trustpolicy.json` file is located. Then enter the following command to create a new role:

```
aws iam create-role --role-name s3ExportRole --assume-role-policy-document
file://trustpolicy.json
```

If the command runs successfully, you will see an output similar to the following:

- At the command line, enter the following command to attach the policy that you created in the previous section to the role that you created:

```
aws iam attach-role-policy --policy-arn arn:aws:iam::123456789012:policy/
s3ExportPolicy --role-name s3ExportRole
```

In the preceding command, replace *arn:aws:iam::123456789012:policy/s3ExportPolicy* with the ARN of the policy that you created in the previous section.

Retrieving recommendations from Amazon Personalize

You can configure Amazon Pinpoint to retrieve recommendation data from an Amazon Personalize solution that's been deployed as an Amazon Personalize campaign. You can use this data to send personalized recommendations to message recipients based on each recipient's attributes and behavior. To learn more, see [Machine learning models](#) in the *Amazon Pinpoint User Guide*.

Before you can retrieve recommendation data from an Amazon Personalize campaign, you must create an AWS Identity and Access Management (IAM) role that allows Amazon Pinpoint to retrieve the data from the campaign. Amazon Pinpoint can create this role for you automatically when you use the console to set up a recommender model in Amazon Pinpoint. Or, you can create this role manually.

To create the role manually, use the IAM API to complete the following steps:

1. Create an IAM policy that allows an entity (in this case, Amazon Pinpoint) to retrieve recommendation data from an Amazon Personalize campaign.
2. Create an IAM role and attach the IAM policy to it.

This topic explains how to complete these steps by using the AWS Command Line Interface (AWS CLI). It assumes that you've already created the Amazon Personalize solution and deployed it as an Amazon Personalize campaign. For information about creating and deploying a campaign, see [Creating a campaign](#) in the *Amazon Personalize Developer Guide*.

This topic also assumes that you've already installed and configured the AWS CLI. For information about setting up the AWS CLI, see [Get started with the AWS CLI](#) in the *AWS Command Line Interface User Guide*.

Step 1: Create the IAM policy

An IAM policy defines permissions for an entity, such as an identity or resource. To create a role that allows Amazon Pinpoint to retrieve recommendation data from an Amazon Personalize campaign, you must first create an IAM policy for the role. This policy must allow Amazon Pinpoint to do the following:

- Retrieve configuration information for the solution that's deployed by the campaign (`DescribeSolution`).
- Check the status of the campaign (`DescribeCampaign`).
- Retrieve recommendation data from the campaign (`GetRecommendations`).

In the following procedure, the example policy allows this access for a particular Amazon Personalize solution that was deployed by a particular Amazon Personalize campaign.

To create the IAM policy

1. In a text editor, create a new file. Paste the following code into the file:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
```

```

        "Sid": "RetrieveRecommendationsOneCampaign",
        "Effect": "Allow",
        "Action": [
            "personalize:DescribeSolution",
            "personalize:DescribeCampaign",
            "personalize:GetRecommendations"
        ],
        "Resource": [
            "arn:aws:personalize:us-east-1:111122223333:solution/solutionId",
            "arn:aws:personalize:us-east-1:111122223333:campaign/campaignId"
        ]
    }
}
}
}

```

In the preceding example, replace the *italicized* text with your information:

- *region* – The name of the AWS Region that hosts the Amazon Personalize solution and campaign.
 - *accountId* – Your AWS account ID.
 - *solutionId* – The unique resource ID for the Amazon Personalize solution that's deployed by the campaign.
 - *campaignId* – The unique resource ID for the Amazon Personalize campaign from which to retrieve recommendation data.
2. When you finish, save the file as `RetrieveRecommendationsPolicy.json`.
 3. By using the command line interface, navigate to the directory where you saved the `RetrieveRecommendationsPolicy.json` file.
 4. Enter the following command to create a policy and name it `RetrieveRecommendationsPolicy`. To use a different name, change *RetrieveRecommendationsPolicy* to the name that you want.

```
aws iam create-policy --policy-name RetrieveRecommendationsPolicy --policy-document file://RetrieveRecommendationsPolicy.json
```

If the policy was created successfully, you will see an output similar to the following:

```
{
  "Policy": {
    "PolicyName": "RetrieveRecommendationsPolicy",
    "PolicyId": "ANPAJ2YJQRJCG3EXAMPLE",
    "Arn": "arn:aws:iam::123456789012:policy/RetrieveRecommendationsPolicy",
    "Path": "/",
    "DefaultVersionId": "v1",
    "AttachmentCount": 0,
    "PermissionBoundaryUsageCount": 0,
    "IsAttachable": true,
    "CreateDate": "2020-03-04T22:23:15Z",
    "UpdateDate": "2020-03-04T22:23:15Z"
  }
}
```

Note

If you receive a message that your account isn't authorized to perform the `CreatePolicy` operation, you must attach a policy to your user account that lets you create new IAM policies and roles for your account. For more information, see [Adding and removing IAM identity permissions](#) in the *IAM User Guide*.

5. Copy the Amazon Resource Name (ARN) of the policy (`arn:aws:iam::123456789012:policy/RetrieveRecommendationsPolicy` in the preceding example). You need this ARN to create the IAM role in the next section.

Step 2: Create the IAM role

After you create the IAM policy, you can create an IAM role and attach the policy to it.

Each IAM role contains a *trust policy*, which is a set of rules that specifies which entities are allowed to assume the role. In this section, you create a trust policy that allows Amazon Pinpoint to assume the role. Next, you create the role itself. Then, you attach the policy to the role.

To create the IAM role

1. In a text editor, create a new file. Paste the following code into the file:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "pinpoint.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

2. Save the file as `RecommendationsTrustPolicy.json`.
3. By using the command line interface, navigate to the directory where you saved the `RecommendationsTrustPolicy.json` file.
4. Enter the following command to create a new role and name it `PinpointRoleforPersonalize`. To use a different name, change *PinpointRoleforPersonalize* to the name that you want.

```
aws iam create-role --role-name PinpointRoleforPersonalize --assume-role-policy-document file://RecommendationsTrustPolicy.json
```

If the command runs successfully, you will see an output similar to the following:

5. Enter the following command to attach the policy that you created in the previous section to the role that you created:

```
aws iam attach-role-policy --policy-arn arn:aws:iam::123456789012:policy/RetrieveRecommendationsPolicy --role-name PinpointRoleforPersonalize
```

In the preceding command, replace *arn:aws:iam::123456789012:policy/RetrieveRecommendationsPolicy* with the ARN of the policy that you created in the previous section. Also, replace *PinpointRoleforPersonalize* with the name of the role that you specified in step 4, if you specified a different name for the role.

IAM role for streaming events to Kinesis

Amazon Pinpoint can automatically send *app usage data*, or *event data*, from your app to an Amazon Kinesis data stream or Amazon Data Firehose delivery stream in your AWS account. Before Amazon Pinpoint can begin streaming the event data, you must delegate the required permissions to Amazon Pinpoint.

If you use the console to set up event streaming, Amazon Pinpoint will automatically create an AWS Identity and Access Management (IAM) role with the required permissions. For more information, see [Streaming Amazon Pinpoint events to Kinesis](#) in the *Amazon Pinpoint User Guide*.

If you want to create the role manually, attach the following policies to the role:

- A permissions policy that allows Amazon Pinpoint to send event data to your stream.
- A trust policy that allows Amazon Pinpoint to assume the role.

After you create the role, you can configure Amazon Pinpoint to send events to your stream automatically. For more information, see [Streaming Amazon Pinpoint events to Kinesis](#) in the *Amazon Pinpoint Developer Guide*.

Permissions policies

To allow Amazon Pinpoint to send event data to your stream, attach one of the following policies to the role.

Amazon Kinesis Data Streams

The following policy allows Amazon Pinpoint to send event data to a Kinesis stream.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "kinesis:PutRecords",
        "kinesis:DescribeStream"
      ],
      "Effect": "Allow",
      "Resource": [
```

```

        "arn:aws:kinesis:us-east-1:111122223333:stream/stream-name"
    ]
}
}

```

Amazon Data Firehose

The following policy allows Amazon Pinpoint to send event data to a Firehose delivery stream.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": {
    "Effect": "Allow",
    "Action": [
      "firehose:PutRecordBatch",
      "firehose:DescribeDeliveryStream"
    ],
    "Resource": [
      "arn:aws:firehose:us-east-1:111122223333:deliverystream/delivery-stream-
name"
    ]
  }
}

```

Trust policy

To allow Amazon Pinpoint to assume the IAM role and perform the actions allowed by the permissions policy, attach the following trust policy to the role:

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {

```

```
    "Service": "pinpoint.amazonaws.com"
  },
  "Action": "sts:AssumeRole"
}
]
```

Creating the IAM role (AWS CLI)

Complete the following steps to create the IAM role by using the AWS Command Line Interface (AWS CLI). To learn how to create the role by using the Amazon Pinpoint console, see [Streaming Amazon Pinpoint events to Kinesis](#) in the *Amazon Pinpoint User Guide*.

If you haven't installed the AWS CLI, see [Installing the AWS CLI](#) in the *AWS Command Line Interface User Guide*.

To create the IAM role by using the AWS CLI

1. Create a JSON file that contains the trust policy for your role, and save the file locally. You can copy the trust policy provided in this topic.
2. Use the [create-role](#) command to create the role and attach the trust policy:

```
aws iam create-role --role-name PinpointEventStreamRole --assume-role-policy-document file://PinpointEventStreamTrustPolicy.json
```

Following the `file://` prefix, specify the path to the JSON file that contains the trust policy.

After you run this command, the AWS CLI prints the following output in your terminal:

3. Create a JSON file that contains the permissions policy for your role, and save the file locally. You can copy one of the policies provided in the [Permissions policies](#) section of this topic.
4. Use the [put-role-policy](#) command to attach the permissions policy to the role:

```
aws iam put-role-policy --role-name PinpointEventStreamRole --policy-name PinpointEventStreamPermissionsPolicy --policy-document file://PinpointEventStreamPermissionsPolicy.json
```

Following the `file://` prefix, specify the path to the JSON file that contains the permissions policy.

IAM role for streaming email events to Firehose

In the Amazon Pinpoint Email API, you can create *configuration sets* that specify how to handle certain email events. For example, you can create a configuration set that sends delivery notifications to a specific *event destination*, such as an Amazon SNS topic or an Amazon Data Firehose delivery stream. When you send email through the Amazon Pinpoint Email API using that configuration set, Amazon Pinpoint sends information about email-related events to the event destination that you specified in the configuration set.

The Amazon Pinpoint Email API can deliver information about the following types of email events to the event destinations that you specify:

- **Sends** – The call to Amazon Pinpoint was successful, and Amazon Pinpoint attempted to deliver the email.
- **Deliveries** – Amazon Pinpoint successfully delivered the email to the recipient's mail server.
- **Rejections** – Amazon Pinpoint accepted the email, determined that it contained malware, and rejected it. Amazon Pinpoint didn't attempt to deliver the email to the recipient's mail server.
- **Rendering Failures** – The email wasn't sent because of a template rendering issue. This event type only occurs when you send an email that includes substitution tags. This event type can occur when substitution values are missing. It can also occur when there's a mismatch between the substitution tags that you used in the email and the substitution data that you provided.

Note

If you use substitution tags in the emails that you send by using the Amazon Pinpoint Email API, you should always create a configuration set that records Rendering Failure events.

- **Bounces** – The recipient's mail server permanently rejected the email.
- **Complaints** – The email was successfully delivered to the recipient, but the recipient used the "Report Spam" (or equivalent) feature of their email client to report the message.
- **Opens** – The recipient received the message and opened it in their email client.
- **Clicks** – The recipient clicked one or more links that were contained in the email.

Note

Every time a recipient opens or clicks an email, Amazon Pinpoint generates unique open or click events, respectively. In other words, if a specific recipient opens a message five times, Amazon Pinpoint reports five separate Open events.

If you want to send data about these events to a Firehose stream, you must create an IAM role that has the appropriate permissions. The role must use the following policies:

- A trust policy that allows Amazon Pinpoint to assume the role.
- A permissions policy that allows the Amazon Pinpoint Email API to send email delivery and response records to your stream.

After you create the role, you can configure Amazon Pinpoint to send events to your stream automatically. For more information, see [Streaming Amazon Pinpoint events to Kinesis](#) in the *Amazon Pinpoint Developer Guide*.

Trust policy

To allow the Amazon Pinpoint Email API to assume the IAM role and perform the actions allowed by the permissions policy, attach the following trust policy to the role:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "",
      "Effect": "Allow",
      "Principal": {
        "Service": "ses.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "sts:ExternalId": "accountId"
        }
      }
    }
  ]
}
```

```

    }
  }
]
}

```

In the preceding example, replace *accountId* with the ID of your AWS account.

Permissions policy

To allow the Amazon Pinpoint Email API to send email event data to a Firehose delivery stream, attach the following permissions policy to a role.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "firehose:PutRecordBatch",
        "firehose:DescribeDeliveryStream"
      ],
      "Resource": [
        "arn:aws:firehose:us-east-1:111122223333:deliverystream/deliveryStreamName"
      ]
    }
  ]
}

```

In the preceding example, replace *region* with the name of the AWS Region in which you created the delivery stream. Replace *accountId* with the ID of your AWS account. Finally, replace *deliveryStreamName* with the name of the delivery stream.

Creating the IAM role (AWS CLI)

Complete the following steps to create the IAM role by using the AWS Command Line Interface (AWS CLI). For information about installing and configuring the AWS CLI, see [Installing the AWS CLI](#) in the *AWS Command Line Interface User Guide*.

To create the IAM role by using the AWS CLI

1. Create a JSON file that contains the trust policy for your role, and then save the file locally. You can copy the [trust policy](#) that's provided earlier in this topic.
2. Use the [create-role](#) command to create the role and attach the trust policy:

```
aws iam create-role --role-name PinpointEventStreamRole \  
--assume-role-policy-document file://PinpointEventStreamTrustPolicy.json
```

In the preceding example, replace *PinpointEventStreamTrustPolicy.json* with the full path to the file that contains the trust policy.

After you run this command, the AWS CLI returns the following output:

3. Create a JSON file that contains the permissions policy for your role, and then save the file locally. You can copy the [permissions policy](#) that's provided earlier in this topic.
4. Use the [put-role-policy](#) command to attach the permissions policy to the role:

```
aws iam put-role-policy \  
--role-name PinpointEventStreamRole \  
--policy-name PinpointEventStreamPermissionsPolicy \  
--policy-document file://PinpointEventStreamPermissionsPolicy.json
```

In the preceding example, replace *PinpointEventStreamPermissionsPolicy.json* with the full path to the file that contains the permissions policy.

Troubleshooting Amazon Pinpoint identity and access management

Use the following information to diagnose and fix common issues that you might encounter when working with Amazon Pinpoint and IAM.

Topics

- [I'm not authorized to perform an action in Amazon Pinpoint](#)
- [I'm not authorized to perform iam:PassRole](#)
- [I want to allow people outside my AWS account to access my Amazon Pinpoint resources](#)

I'm not authorized to perform an action in Amazon Pinpoint

If the AWS Management Console tells you that you're not authorized to perform an action, then you must contact your administrator for assistance. Your administrator is the person who provided you with your user name and password.

The following example error occurs when the `mateojackson` IAM user tries to use the console to view details about a project but doesn't have `mobiletargeting:GetApp` permissions.

```
User: arn:aws:iam::123456789012:user/mateojackson is not authorized to perform:
mobiletargeting:GetApp on resource: my-example-project
```

In this case, Mateo asks his administrator to update his policies to allow him to access the `my-example-project` resource using the `mobiletargeting:GetApp` action.

I'm 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 Amazon Pinpoint.

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 Amazon Pinpoint. 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 my AWS account to access my Amazon Pinpoint 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 Amazon Pinpoint supports these features, see [How Amazon Pinpoint works with IAM](#).
- To learn how to provide access to your resources across AWS accounts that you own, see [Providing access to an IAM user in another AWS account that you own](#) in the *IAM User Guide*.
- To learn how to provide access to your resources to third-party AWS accounts, see [Providing access to AWS accounts owned by third parties](#) in the *IAM User Guide*.
- To learn how to provide access through identity federation, see [Providing access to externally authenticated users \(identity federation\)](#) in the *IAM User Guide*.
- To learn the difference between using roles and resource-based policies for cross-account access, see [Cross account resource access in IAM](#) in the *IAM User Guide*.

Event logging and monitoring in Amazon Pinpoint

Logging and monitoring are an important part of maintaining the reliability, availability, and performance of your Amazon Pinpoint projects and other types of Amazon Pinpoint resources. You can log and collect monitoring data from all parts of your Amazon Pinpoint projects and resources to streamline the process of debugging a multipoint failure, if one occurs. AWS provides several tools that can help you log and collect this data and respond to potential incidents:

AWS CloudTrail

Amazon Pinpoint integrates with AWS CloudTrail, which is a service that provides a record of actions that were taken in Amazon Pinpoint by a user, a role, or another AWS service. This includes actions from the Amazon Pinpoint console and programmatic calls to Amazon Pinpoint API operations. By using the information collected by CloudTrail, you can determine which requests were made to Amazon Pinpoint. For each request, you can identify when it was made, the IP address from which it was made, who made it, and additional details. For

more information, see [Logging Amazon Pinpoint API calls with AWS CloudTrail](#) in the *Amazon Pinpoint Developer Guide*.

Amazon CloudWatch

You can use Amazon CloudWatch to collect, view, and analyze several important metrics related to your Amazon Pinpoint account and projects. You can also use CloudWatch to create alarms that notify you if the value for a metric meets certain conditions and is within or exceeds a threshold that you define. If you create an alarm, CloudWatch sends a notification to an Amazon Simple Notification Service (Amazon SNS) topic that you specify. For more information, see [Monitoring Amazon Pinpoint with Amazon CloudWatch](#) in the *Amazon Pinpoint User Guide*.

AWS Health Dashboards

By using AWS Health dashboards, you can check and monitor the status of your Amazon Pinpoint environment. To check the status of the Amazon Pinpoint service overall, use the AWS Service Health Dashboard. To check, monitor, and view historical data about any events or issues that might affect your AWS environment more specifically, use the AWS Personal Health Dashboard. To learn more about these dashboards, see the [AWS Health User Guide](#).

AWS Trusted Advisor

AWS Trusted Advisor inspects your AWS environment and provides recommendations for opportunities to address security gaps, improve system availability and performance, and save money. All AWS customers have access to a core set of Trusted Advisor checks. Customers who have a Business or Enterprise Support plan have access to additional Trusted Advisor checks.

Many of these checks can help you assess the security posture of your Amazon Pinpoint resources as part of your AWS account overall. For example, the core set of Trusted Advisor checks includes the following:

- Logging configurations for your AWS account, for each supported AWS Region.
- Accessing permissions for your Amazon Simple Storage Service (Amazon S3) buckets, which might contain files that you import into Amazon Pinpoint to build segments.
- Using AWS Identity and Access Management (IAM) users, groups, and roles to control access to Amazon Pinpoint resources.
- Identifying IAM configurations and policy settings that might compromise the security of your AWS environment and Amazon Pinpoint resources.

For more information, see [AWS Trusted Advisor](#) in the *Support User Guide*.

Compliance validation for Amazon Pinpoint

Third-party auditors assess the security and compliance of Amazon Pinpoint as part of multiple AWS compliance programs. These include AWS System and Organization Controls (SOC), FedRAMP, HIPAA, ISO/IEC 27001:2013 for security management controls, ISO/IEC 27017:2015 for cloud-specific controls, ISO/IEC 27018:2014 for personal data protection, ISO/IEC 9001:2015 for quality management systems, and others.

For a list of AWS services that are in scope for specific compliance programs, see [AWS services in scope by compliance program](#). For general information, see [AWS compliance programs](#).

You can download third-party audit reports by using AWS Artifact. For more information, see [Downloading reports in AWS Artifact](#).

Your compliance responsibility when using Amazon Pinpoint 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:

- [Quick Starts](#) – These deployment guides discuss architectural considerations and provide steps for deploying security- and compliance-focused baseline environments on AWS.
- [Architecting for HIPAA security and compliance on Amazon Web Services whitepaper](#) – This whitepaper describes how companies can use AWS to create HIPAA-compliant applications.
- [AWS compliance resources](#) – This collection of workbooks and guides might apply to your industry and location.
- [Evaluating resources with AWS Config Rules](#) in the *AWS Config Developer Guide* – The AWS Config service assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.
- [AWS Security Hub CSPM](#) – This AWS service provides a comprehensive view of your security state within AWS and helps you check your compliance with security standards and best practices.

Amazon Pinpoint is an AWS HIPAA-eligible service when you use the proper communication channels. If you want to use Amazon Pinpoint to run workloads containing Protected Health Information (PHI), as defined by HIPAA and associated legislation and regulations, you should use the email channel, push notification channel, or SMS channel to send messages that contain PHI. If you use the SMS channel to send messages that contain PHI, you should send those messages from a dedicated [short code](#) that you requested for your AWS account for the explicit purpose of

sending messages that will or may contain PHI. The voice channel is not AWS HIPAA-eligible. Don't use the voice channel to send messages that contain PHI.

Infrastructure security in Amazon Pinpoint

As a managed service, Amazon Pinpoint is protected by AWS global network security. For information about AWS security services and how AWS protects infrastructure, see [AWS Cloud Security](#). To design your AWS environment using the best practices for infrastructure security, see [Infrastructure Protection](#) in *Security Pillar AWS Well-Architected Framework*.

You use AWS published API calls to access Amazon Pinpoint through the network. Clients must support the following:

- Transport Layer Security (TLS). We require TLS 1.2 and recommend TLS 1.3.
- Cipher suites with perfect forward secrecy (PFS) such as DHE (Ephemeral Diffie-Hellman) or ECDHE (Elliptic Curve Ephemeral Diffie-Hellman). Most modern systems such as Java 7 and later support these modes.

Although you can make these API calls from any network location, Amazon Pinpoint supports resource-based access policies. These policies can include restrictions based on source IP address. To learn more about this type of policy, see [Managing access using policies](#).

In addition, you can configure and use various AWS security features to control access to Amazon Pinpoint resources from any mobile or web apps that you integrate with Amazon Pinpoint. This includes restrictions on API calls for tasks such as adding endpoints, updating endpoint data, submitting event data, and reporting usage data.

To use these features, we recommend that you use the AWS Mobile SDKs or AWS Amplify JavaScript libraries to integrate mobile and web apps with Amazon Pinpoint. For Android or iOS apps, we recommend that you use the AWS Mobile SDK for Android or the AWS Mobile SDK for iOS, respectively. For JavaScript-based mobile or web apps, we recommend that you use the AWS Amplify JavaScript Library for the web or the AWS Amplify JavaScript Library for React Native. To learn more about these resources, see [Getting started with the AWS mobile SDKs](#), and [Getting started with the AWS Amplify library for react native](#).

Security best practices for Amazon Pinpoint

The following best practices are general guidelines and don't represent a complete security solution. Because these best practices might not be appropriate or sufficient for your environment, treat them as helpful considerations rather than prescriptions.

AWS Identity and Access Management (IAM) is a service that you can use to control AWS services, including Amazon Pinpoint. Consider the following best practices as you build your Amazon Pinpoint architecture:

- Use IAM accounts to control access to Amazon Pinpoint API operations, especially operations that create, modify, or delete Amazon Pinpoint resources. For the Amazon Pinpoint API, these resources include projects, campaigns, and journeys. For the Amazon Pinpoint SMS and Voice API, these resources include phone numbers, pools, and configuration sets.
- Create an individual IAM user for each person who manages Amazon Pinpoint resources, including yourself. Don't use AWS account root users to manage Amazon Pinpoint resources.
- Grant users the minimum set of permissions required to perform his or her duties.
- Use IAM groups to effectively manage permissions for multiple users.
- Rotate your IAM credentials regularly.

For more information about Amazon Pinpoint security, see [Security in Amazon Pinpoint](#). For more information about IAM, see [AWS Identity and Access Management](#). For information on IAM best practices, see [Security best practices in IAM](#).

Document history for the *Amazon Pinpoint Resilient Architecture Guide*

The following table describes important changes in each release of the *Amazon Pinpoint Resilient Architecture Guide*. To receive notifications when this guide is updated, you can subscribe to the RSS feed.

- **Latest documentation update:** October 31, 2022

Change	Description	Date
End of support notice	End of support notice: On October 30, 2026, AWS will end support for Amazon Pinpoint. After October 30, 2026, you will no longer be able to access the Amazon Pinpoint console or Amazon Pinpoint resources (endpoints, segments, campaigns, journeys, and analytics). For more information, see Amazon Pinpoint end of support . Note: APIs related to SMS, voice, mobile push, OTP, and phone number validate are not impacted by this change and are supported by AWS End User Messaging.	May 20, 2025
Encryption in transit	Starting 2023-03-22 Amazon Pinpoint will no longer support TLS 1.0 but you can still use TLS 1.2 or later.	March 20, 2023

For more information, see [Encryption in transit](#).

[Initial release](#)

The initial release of the *Amazon Pinpoint Resilient Architecture Guide* contains information, best practices , and example architectures related to building a resilient , multi-Region architecture with Amazon Pinpoint.

October 31, 2022