



Lake Partner Onboarding Guide

Amazon CloudTrail



Version 1.0

Amazon CloudTrail: Lake Partner Onboarding Guide

Table of Contents

Onboarding partner events to Amazon CloudTrail Lake	1
How integrations with CloudTrail Lake add value	1
Terminology	2
How partner integration works	3
Onboard to Amazon CloudTrail Lake	3
Prerequisites	4
Step 1: Partner registration	4
Step 2: Build the integration	4
Best practices and quotas	5
CloudTrail Lake event schema	6
Learn more about CloudTrail Lake	14

Onboarding partner events to Amazon CloudTrail Lake

Amazon CloudTrail Lake logs activities across all Amazon Web Services accounts and Regions, and across a customer's entire IT infrastructure. Customers can configure CloudTrail Lake to log events from any source, immutably store the events for auditing and compliance, and use standard, SQL-based queries to filter and analyze their event logs. CloudTrail Lake accepts activity logs from Amazon Partner Network partner solutions, offering customers a comprehensive view of their activity information in the CloudTrail Lake console, or by using API commands.

This guide is for Amazon Partners who want to explore creating an integration with CloudTrail Lake, or want to know how to onboard their applications and solutions to CloudTrail Lake, and let their customers integrate their activity events into CloudTrail Lake.

Topics

- [How integrations with CloudTrail Lake add value](#)
- [Terminology](#)
- [How partner integration works](#)
- [Onboard to Amazon CloudTrail Lake](#)
- [Understanding the CloudTrail Lake event schema](#)
- [Learn more about CloudTrail Lake](#)

How integrations with CloudTrail Lake add value

As an Amazon Partner, an integration with CloudTrail Lake can add value for you in the following ways:

- **Modern, consolidated solution for audit logging for your customers:** Today, audit and security professionals get trusted records of Amazon activity from CloudTrail. CloudTrail users need a similar experience for other application audit information, regardless of the source. Partner integrations centralize audit logging, and extend CloudTrail benefits, such as immutable storage for 7 years and a query interface for analysis, to partner solutions, simplifying audit and compliance processes for our common customers.
- **Discovery for partners:** CloudTrail promotes partners with integrations in CloudTrail console, including links to partner Amazon Web Services Marketplace listings.

Terminology

The following terms are helpful in understanding how a CloudTrail Lake integration works.

Event data store

CloudTrail Lake lets you run SQL-based queries on your events. Events are aggregated into event data stores, which are immutable collections of events based on criteria that you select by applying [advanced event selectors](#). You can keep the event data in an event data store for up to 3,653 days (about 10 years) if you choose the **One-year extendable retention pricing** option, or up to 2,557 days (about 7 years) if you choose the **Seven-year retention pricing** option. The selectors that you apply to an event data store control which events persist and are available for you to query. For non-Amazon sources, including partner sources, customers create an event data store to log activity events using the CloudTrail console or API. The event type in the console is **Events from integrations**. In the API, the eventCategory value is **ActivityAuditLog**.

Channel

A partner-specific resource that Amazon customers create as part of the integration process in CloudTrail Lake. Channels let customers map event sources to destinations. Channels for onboarded partner events have the partner solution set as the source, and event data stores to which customers want to deliver partner events set as destinations. To finish the integration process, customers provide the partner with an Amazon Resource Name (ARN) of the channel. The partner solution uses the channel to send events to CloudTrail Lake.

Resource policy

A permissions policy that is attached to the channel resource and identifies who has access to the channel.

Direct integration

CloudTrail supports two integration types: direct and solution. With a direct integration, the partner calls the PutAuditEvents API to deliver events to the event data store for the customer's Amazon account.

Solution integration

CloudTrail supports two integration types: direct and solution. With a solution integration, the application runs in the customer's Amazon account and the application calls the

PutAuditEvents API to deliver events to the event data store for the customer's Amazon account.

How partner integration works

The following diagram shows how an Amazon customer configures event integration with an onboarded partner. The diagram assumes that the person who is responsible for managing the Amazon account also manages the partner application. The process is described following the diagram.

1. The Amazon customer [creates an event data store](#).
2. The Amazon customer starts partner integration in the CloudTrail Lake integration page of the Amazon Web Services Management Console, and finishes the workflow. The workflow creates a channel for the partner and attaches a resource policy to the channel. A channel ARN is a unique connection between a partner and an Amazon customer's account.
3. The customer provides the partner application with the channel ARN.
4. The customer performs an auditable activity that generates an event in the partner application.
5. The partner sends the audit event to CloudTrail Lake by calling the [PutAuditEvents API](#), and using it to pass the eventData content from the customer's activity, the channel ARN, and the external ID (if included in the resource policy).
6. CloudTrail Lake checks the resource policy to verify that the partner's permissions are valid. If the partner's permissions are valid, CloudTrail Lake ingests the activity events.

Onboard to Amazon CloudTrail Lake

This section describes the prerequisites and steps to onboard your partner application to CloudTrail Lake.

Topics

- [Prerequisites](#)
- [Step 1: Partner registration](#)
- [Step 2: Build the integration](#)
- [Best practices and quotas](#)

Prerequisites

The following are requirements for performing tasks in this guide.

- Amazon provides tiers (Select, Advanced, Premier) to recognize organizations that have proven technical expertise and demonstrated customer experience. You must be at least an [Amazon Select Tier Partner](#). To become an Amazon Partner, you must first meet all [requirements](#) for the tier.

For more information about how to become an Amazon Select Tier partner, see [Become an Amazon Partner](#).

Step 1: Partner registration

To get started, register as an Amazon Partner in the Amazon Partner Network.

Be sure to meet the requirements of partner intake forms. The partner CloudTrail Lake intake forms collect information that the Amazon Partner Network uses to create your partner product profile. This profile gives the CloudTrail team information that we add to your partner provider description that is displayed in the CloudTrail console. Your profile also includes information that CloudTrail uses to confirm the integrity of the event source as CloudTrail Lake receives events from a partner application.

1. Get started by [joining the Amazon Partner Network](#), and informing your Amazon Partner Network team that you want to become a partner with CloudTrail Lake.
2. Get onboarding materials—including partner onboarding forms and the CloudTrail event schema—from the Amazon Partner Network team.
3. Complete the partner onboarding forms, and share the completed forms with your Amazon Partner Network team. You might not yet have all required details. If you have questions, contact your Amazon Partner Network team.

Step 2: Build the integration

Build the integration that is required to send event logs to CloudTrail Lake.

1. Review the [CloudTrail integration event schema](#) in this guide. The CloudTrail event schema provides a consistent way to log activity events for audit needs. This eliminates the need for

time-consuming data standardization efforts before a cross-source analysis. CloudTrail Lake cannot accept events that do not follow the prescribed schema.

2. Determine the events that you want to send. CloudTrail Lake only accepts activity events, or events that help customers understand who did what, and when. Typically, partners have existing mechanisms to provide their customers access to activity logs. The schema mapping exercise helps you exclude non-activity events. Contact your Amazon Partner Network team if you need help narrowing down event types.
3. Build your integration architecture to send activity events to CloudTrail Lake. This includes offering a setup framework (GUI is preferred) and documentation for customers to enable your partner application to send events to CloudTrail Lake. A partner customer must share a CloudTrail channel Amazon Resource Number (ARN) with the partner as part of the integration process.
 - a. To send events to CloudTrail Lake, the partner calls the [PutAuditEvents API](#), specifying the channel ARN provided by the customer. If the channel's resource policy includes an external ID, you must also pass the external ID when you call PutAuditEvents.
 - b. The partner checks transfer results for failures, and tries to resend failed events by calling the PutAuditEvents API again.

Best practices and quotas

As you integrate partner solution events, be aware of the following best practices, quotas, and limitations.

- **Schema mapping:** Be sure that you have the key required fields included in the eventData block. Missing required fields results in errors. For information about required fields, see [Understanding the CloudTrail Lake event schema](#)

You can add event fields that do not map to the schema to the additionalEventData field. Some partners use this field to include the entire, raw event.

- **Batching events:** When you call the PutAuditEvents API, you can batch up to 100 events in a single API call, as long as each event is not greater than 256 kB in size, and the total size of all events is less than 1 MB. For more information about quotas in CloudTrail, see [Quotas in Amazon CloudTrail](#) in the *Amazon CloudTrail User Guide*.

Understanding the CloudTrail Lake event schema

The tables in this section describe the required and optional schema elements that match those in CloudTrail event records. The contents of eventData are provided by customer events; other fields are provided by CloudTrail after customer events are ingested.

- [Fields that are provided by CloudTrail after ingestion](#)
- [Fields that are provided by your events](#)

The following fields are provided by CloudTrail after ingestion:

Field name	Input type	Requirement	Description
eventVersion	string	Required	The event version.
eventCategory	string	Required	The event category. For non-Amazon events, the value is ActivityAuditLog .
eventType	string	Required	The event type. For non-Amazon events, the valid value is ActivityLog .
eventID	string	Required	A unique ID for an event.
eventTime	string	Required	Event timestamp , in yyyy-MM-DTHH:mm:ss format, in Universal Coordinated Time (UTC).

Field name	Input type	Requirement	Description
awsRegion	string	Required	The Amazon Web Services Region where the PutAuditEvents call was made.
recipientAccountId	string	Required	Represents the account ID that received this event. CloudTrail populates this field by calculating it from event payload.
addendum	-	Optional	Shows information about why event processing was delayed. If information was missing from an existing event, the addendum block includes the missing information and a reason for why it was missing.
• reason	string	Optional	The reason that the event or some of its contents were missing.

Field name	Input type	Requirement	Description
• updatedFields	string	Optional	The event record fields that are updated by the addendum. This is only provided if the reason is UPDATED_DATA .
• originalUID	string	Optional	The original event UID from the source. This is only provided if the reason is UPDATED_DATA .
• originalEventID	string	Optional	The original event ID. This is only provided if the reason is UPDATED_DATA .
metadata	-	Required	Information about the channel that the event used.
• ingestionTime	string	Required	The timestamp when the event was processed, in yyyy-MM-DDTHH:mm:ss format, in Universal Coordinated Time (UTC).
• channelARN	string	Required	The ARN of the channel that the event used.

The following fields are provided by customer events:

Field name	Input type	Requirement	Description
eventData	-	Required	The audit data sent to CloudTrail in a PutAuditEvents call.
• version	string	Required	The version of the event from its source. Length constraints: Maximum length of 256.
• userIdentity	-	Required	Information about the user who made a request.
• • type	string	Required	The type of user identity. Length constraints: Maximum length of 128.
• • principalId	string	Required	A unique identifier for the actor of the event. Length constraints: Maximum length of 1024.
• • details	JSON object	Optional	Additional information about the identity.

Field name	Input type	Requirement	Description
• userAgent	string	Optional	<p>The agent through which the request was made.</p> <p>Length constraints: Maximum length of 1024.</p>
• eventSource	string	Required	<p>This is the partner event source, or the custom application about which events are logged.</p> <p>Length constraints: Maximum length of 1024.</p>
• eventName	string	Required	<p>The requested action, one of the actions in the API for the source service or application.</p> <p>Length constraints: Maximum length of 1024.</p>
• eventTime	string	Required	Event timestamp, in yyyy-MM-DTHH:mm:ss format, in Universal Coordinated Time (UTC).

Field name	Input type	Requirement	Description
• UID	string	Required	<p>The UID value that identifies the request. The service or application that is called generates this value.</p> <p>Length constraints: Maximum length of 1024.</p>
• requestParameters	JSON object	Optional	<p>The parameters, if any, that were sent with the request. This field has a maximum size of 100 kB, and content exceeding the limit is rejected.</p>
• responseElements	JSON object	Optional	<p>The response element for actions that make changes (create, update, or delete actions). This field has a maximum size of 100 kB, and content exceeding the limit is rejected.</p>
• errorCode	string	Optional	<p>A string representing an error for the event.</p> <p>Length constraints: Maximum length of 256.</p>

Field name	Input type	Requirement	Description
• errorMessage	string	Optional	<p>The description of the error.</p> <p>Length constraints: Maximum length of 256.</p>
• sourceIPAddress	string	Optional	<p>The IP address from which the request was made. Both IPv4 and IPv6 addresses are accepted.</p>
• recipientAccountId	string	Required	<p>Represents the account ID that received this event. The account ID must be the same as the Amazon account ID that owns the channel.</p>
• additionalEventData	JSON object	Optional	<p>Additional data about the event that was not part of the request or response. This field has a maximum size of 28 kB, and content exceeding that limit is rejected.</p>

The following example shows the hierarchy of schema elements that match those in CloudTrail event records.

```
{  
  "eventVersion": String,  
  "eventCategory": String,  
  "eventType": String,  
  "eventID": String,  
  "eventTime": String,  
  "awsRegion": String,  
  "recipientAccountId": String,  
  "addendum": {  
    "reason": String,  
    "updatedFields": String,  
    "originalUID": String,  
    "originalEventID": String  
  },  
  "metadata" : {  
    "ingestionTime": String,  
    "channelARN": String  
  },  
  "eventData": {  
    "version": String,  
    "userIdentity": {  
      "type": String,  
      "principalId": String,  
      "details": {  
        JSON  
      }  
    },  
    "userAgent": String,  
    "eventSource": String,  
    "eventName": String,  
    "eventTime": String,  
    "UID": String,  
    "requestParameters": {  
      JSON  
    },  
    "responseElements": {  
      JSON  
    },  
    "errorCode": String,  
    "errorMessage": String,  
    "sourceIPAddress": String,  
    "recipientAccountId": String,  
    "additionalEventData": {  
      JSON  
    }  
  }  
}
```

```
        JSON
    }
}
}
```

Learn more about CloudTrail Lake

The following resources can help you get a better understanding of what CloudTrail Lake is and how Amazon customers use it. We encourage you to try CloudTrail in one of your Amazon Web Services accounts, and get more experience using the service.

- [Modernize Your Audit Log Management Using CloudTrail Lake](#) (YouTube video)
- [Log Activity Events from Non-Amazon Sources in Amazon CloudTrail Lake](#) (YouTube video)
- [Analyze Activity Logs with Amazon CloudTrail Lake and Amazon Athena](#) (YouTube video)
- [Get visibility into the activity logs for your workforce and customer identities](#) (Amazon blog)
- [Using Amazon CloudTrail Lake to identify older TLS connections to Amazon service endpoints](#) (Amazon blog)
- [How Arctic Wolf uses Amazon CloudTrail Lake to Simplify Security and Operations](#) (Amazon blog)
- [Working with CloudTrail Lake in the Amazon CloudTrail User Guide](#)
- [Amazon CloudTrail API Reference](#)
- [Amazon CloudTrail Data API Reference](#)