

Choosing an AWS application integration service



Choosing an AWS application integration service: AWS Decision Guide

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

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

Table of Contents

Decision guide

Introduction

Understand

Define

Consider

Choose

Use

Explore

i

1

2

3

5

7

10

18

Document history

20

Choosing an AWS application integration service

Taking the first step

Purpose	Help determine which AWS application integration services are the best fit for your workloads.
Last updated	April 16, 2025
Covered services	<ul style="list-style-type: none">• Amazon SNS• Amazon SQS• Amazon EventBridge• Amazon MQ• Amazon Kinesis Data Streams• Amazon MSK• Amazon MWAA• AWS Step Functions

Introduction

Application integration is a suite of services that enables communication between decoupled components within microservices, distributed systems, and serverless applications. Amazon Web Services (AWS) offers more than half a dozen application integration services to support a diverse set of workloads running in the cloud.

Choosing an integration service that is the best fit for your organization and workloads can become difficult. This decision guide will help you ask the right questions to discover your requirements and provides clear guidance on how to evaluate and choose the right integration services for your workloads.

[This 8½-minute clip is from a one hour recording of a presentation by AWS director of enterprise strategy Gregor Hohpe at AWS re:Invent 2022. It provides an overview of available AWS application integration services.](#)

Understand

One of the key benefits of modernization is the ability to shift operational responsibilities, allowing you to free up resources to do more value-add and innovation-led activities.

There is a spectrum of shared responsibility options across different levels of modernization, ranging from hosting your message broker on Amazon Elastic Compute Cloud (Amazon EC2)—where you are managing scaling, security configurations, provisioning, patching, and more—to serverless offerings where all of the underlying infrastructure is managed.



As you start to explore and understand your criteria, environment, and the suite of integration services that AWS offers, we recommend that you review some best practices. These best practices are applicable regardless of which service (or suite of services) you choose.

Understand integration in your environment

It's common that some organizations spend more time than they would like on maintaining open-source integrations. We recommend that you consider community sources, and/or backing from enterprises or foundations when making these investments. An investment in these projects is not

just financial, but also an investment in knowledge capital and potentially technical debt, as these components and associated integrations will typically need updating. For more information, see the [AWS Open Source blog](#).

Understand your architecture characteristics

The ability to support a wide range of architectures is important. We recommend that you use the [AWS Well-Architected Framework](#) as a guide to help you understand the decisions you make when building architectures on AWS. Additionally, using the Well-Architected Framework allows you to learn architectural best practices for designing and operating reliable, scalable, secure, efficient, and cost-effective systems in the Cloud.

Use a combination of integration services

If you are using purpose-built services, a combination of services might be the best fit for your use case. The following lists out a few common ways AWS customers are using a combination of services.

- Routing Amazon EventBridge or Amazon Simple Notification Service (Amazon SNS) events to an Amazon Simple Queue Service (Amazon SQS) queue as a buffer for downstream consumers.
- Pulling events directly from a stream (Amazon Kinesis Data Streams or Amazon Managed Streaming for Apache Kafka (Amazon MSK)) or a queue (Amazon SQS or Amazon MQ) with EventBridge Pipes and sending events to an EventBridge bus to push out to consumers.
- Routing EventBridge or Amazon SNS events to a Kinesis Data Streams or Amazon MSK for gathering and viewing analytics.

Define

Once you have a clearer picture of your criteria, environment, strategic direction, and available services (including both deployment hosted and managed modalities), you need to identify your integration requirements. You might already know some of your requirements if you are migrating to an existing integration platform or message broker. However, you need to establish how these requirements would change if you move to a cloud environment, if at all.

Messaging or streaming platforms

These platforms are expected to fulfill a certain business functionality. Use the following example use cases when considering which functionalities you will need.

Example 1:

Consider an insurance company which receives different claims as messages for different claim types (auto, home, or life) with different business rules. It might mean that the message consumer should have the functionality to route claims to a different destination based on header properties in the message.

Example 2:

Consider an airline in which a flight status update needs to notify all connected systems such as baggage or gate operations using a protocol such as Advanced Messaging Queuing Protocol (AMQP). The big question with functional and business use case primitives is what constitutes a best fit messaging platform. We have multiple choices that can determine the suitability of the platform based on the use case.

- **Market adoption:** This platform is widely adopted by a huge customer community and a good enough fit for most use cases. It is tried and tested with a vibrant support community for any issues that might be encountered. It is a low-risk decision with enough training available for the development resources.
- **Best fit for use case:** These platforms will be tailored for specific industry use cases such as airlines, logistics, or healthcare. They might be the best fit for those use cases with ready-made templates available for adoption. These platforms can be easy to get started but can lack the level of adoption in the market as well as flexibility. Adopting this type of platform might require extensive time and resources for validation and building in house expertise.
- **Modern:** These platforms are built with the next generation architecture to address cloud-scale deployments, multi-tenancy, disaster recovery, and serverless type of pricing. Using this type of platform might require some refactoring of workloads for long-term viability. It uses a cloud-native platform and focuses on using the well-architected principles of modern applications.

Example 3:

If the messaging platform is part of a bigger loan processing workflow which needs to be multi-region, the messaging platform also needs to support the same business requirements. If the business needs the ability to recover and rollback to a previous state in rainy day situations, the underlying messaging or streaming platform also needs capability to have some snapshotting or replay capability to recreate the state of the system.

The integration platform you choose should facilitate asynchronous processing of loan applications or act as the store and forward channel for a multi-step media processing workflow. The criticality of the business process would determine the capabilities needed from the messaging or streaming platform.

Consider

When considering a major application integration architecture in the cloud, there are different ways to determine the functional requirements for each of the integration points.

The following is some of the criteria to consider when choosing an application integration service.

Managed service and operation overhead

Consider moving to the cloud to reduce operational cost by standardizing on managed services that shift the operational burden to AWS. Higher levels of abstraction allow developers and operators to focus on their own unique value-add activities, instead of undifferentiated tasks.

Open source

Consider standardization on [open source technologies](#). Open source can enable an organization to find the right skills—and avoid some risk around lock-in.

Making the wrong choices in an open-source ecosystem can lead to being locked into abstractions and homegrown integrations. Additionally, the responsibility for making different open-source components work together often sits with the organization making the choice. This can lead to organizations spending significant time maintaining open-source integrations.

Workload characteristics

When choosing the right integration service, it is important to understand the characteristics of the messages that need to be sent between the applications. Key characteristics like the message format, size, retention and priority can drive the decision of the integration service.

Some integration services are better suited for small text based messages whereas some are designed to support multiple formats such as text and binary and offer bigger message sizes. The need to have a replay capability can also be an important factor alongside ordering of messages in some scenarios.

For example, message ordering can be implemented by using the FIFO functionality offered by Amazon SNS and Amazon SQS. There is also a consideration of having a pull or push based architecture, such as EventBridge or Amazon SNS invoking a Lambda function asynchronously.

A pull based architecture could use services such as Amazon SQS or Kinesis Data Streams, where messages are stored on a queue or a stream and then can be retrieved by a consuming system. Messaging services like Amazon MQ offer capabilities around bigger message payloads and have unlimited retention. However, they do not offer replay capability.

Rapid iteration and feature velocity

If your primary focus is building and iterating quickly, serverless services might provide the best value. Serverless services let you build applications without managing infrastructure. They provide managed functionality and integrations to reduce time spent writing boiler plate code.

Another benefit of serverless when testing new ideas is that these services offer usage based pricing. Your code only runs when the service is invoked, so an experiment does not require an upfront investment.

Application portability

Many applications use certain protocols - such as Advanced Message Queuing Protocol (AMQP) or MQ Telemetry Transport (MQTT) - to connect to a messaging service. Alternatively, they have some library dependency which uses a certain messaging protocol. Examples of such libraries or frameworks include Spring Boot, Celery, or MassTransit.

You might want to preserve such applications for different reasons. In these cases, the choice of your integration service also depends on the support of the required protocols to have portability with your applications.

Automation portability

You might need to have a service that provides compatibility with your infrastructure and deployment tools - and run the same integration system that you host on-premises (such as Apache ActiveMQ, RabbitMQ, and Apache Kafka).

Managed open source services (such as Amazon MQ and Amazon MSK) provide the benefits of the cloud, while being compatible with many popular deployment tools used for on-premises deployments.

If refactoring the application is an option, you can benefit from using serverless services for providing this capability natively, as well as rich integration with a variety of AWS services.

Organization size and skills

The skills of your organization are a major factor when deciding on the right integration service. If your teams are familiar with a self managed product and it meets your needs, then having a managed service for the same provides the path of least impact. Doing so, can help you apply the best practices for the service and focus on value add activities.

Choose

Now that you know the criteria you will use to evaluate your application integration needs, you are ready to choose which AWS service or services are right for your workloads in your environment.

Service type	When would you use it?	What is it optimized for?	Associated services
Capacity	Use when you need to decouple publishers and subscribers and send events to multiple subscribers simultaneously.	Optimized for asynchronous, loosely coupled communication between publishers and subscribers. Events provide flexibility in message routing and delivery and are well-suited for event-driven architectures where events play a central role in initiating actions or workflows.	Amazon EventBridge Amazon SNS
Messaging	Use when you need either pub/sub messaging to broadcast messages to multiple recipients simultaneously,	Optimized for high-throughput, scalable, and reliable asynchronous pub/sub and point-to-point messaging	Amazon SNS Amazon SQS Amazon MQ

Service type	When would you use it?	What is it optimized for?	Associated services
	or point-to-point messaging when you need reliable and asynchronous communication between components.	between distributed components.	
Streaming	Use streaming services such as Amazon Kinesis Data Streams and Amazon Managed Streaming for Apache Kafka (MSK) in scenarios that involve handling and processing real-time streaming data.	Optimized for ingesting, processing, and analyzing large volumes of real-time streaming data for use cases that require real-time analytics, real-time monitoring, data exploration, and other applications that demand the processing of high-velocity data streams.	Amazon Kinesis Data Streams Amazon MSK

Service type	When would you use it?	What is it optimized for?	Associated services
Workflows	Use when you need to design, coordinate, and manage workflows or sequences of tasks in an organized and scalable manner.	Optimized for use cases such as business process management, application orchestration, data pipeline automation, and microservices coordination. Workflows abstract the underlying infrastructure complexity, allowing you to focus on designing and managing your workflows effectively. They are capable of handling dependencies and sequencing, allowing for parallelism and conditional branching while providing fault tolerance, error handling, and retries to ensure reliable workflow execution.	AWS Step Functions Amazon MWAA

Service type	When would you use it?	What is it optimized for?	Associated services
Scheduling	Use scheduling when you need to automate routine tasks such as data processing, backups, or system health checks. Tasks often need to be run at specific times or intervals, such as every night, hour, or minute.	Optimized for reliable, time-based tasks with built-in retry logic. Suitable for workflows requiring precise scheduling and integration with various AWS services.	Amazon EventBridge

Use

You should now have a clear understanding of what each AWS application integration service does—and which one might be right for you. To explore how to use and learn more about each of the available AWS application integration services—we have provided a pathway to explore how each of the services work. The following section provides links to in-depth documentation, hands-on tutorials, and resources to get you started.

Amazon SNS

- **Filter Messages Published to Topics with Amazon SNS and Amazon SQS**

Learn how to use the message filtering feature of Amazon SNS.

[Get started with the tutorial](#)

- **Amazon SNS - Troubleshooting**

Learn how to view configuration information, monitor processes, and gather diagnostic data about Amazon SNS.

[Explore the guide](#)

- **Build a turn-based game with Amazon DynamoDB and Amazon SNS**

Learn how to build a multiplayer, turn-based game using Amazon DynamoDB and Amazon SNS.

[Get started with the tutorial](#)

- **Building event driven architectures**

Learn how to build on a simple pub/sub implementation using Amazon SNS as our publishing service and Amazon SQS as a subscriber.

[Explore the guide](#)

- **Archiving and replaying messages with Amazon SNS FIFO**

Learn how to archive and replay messages published to Amazon SNS FIFO, which can be useful in failure recovery and state replication scenarios.

[Read the blog post](#)

Amazon SQS

- **Getting started with Amazon SQS**

This guide shows you how to manage queues and messages using the Amazon SQS console.

[Explore the guide](#)

- **Send Fanout Event Notifications**

Learn how to implement a fanout messaging scenario using Amazon SNS and Amazon SQS.

[Get started with the tutorial](#)

- **Orchestrate Queue-based Microservices**

Learn how to design and run a serverless workflow that orchestrates a message queue-based microservice.

[Get started with the tutorial](#)

- **Send Messages Between Distributed Applications**

Use the Amazon SQS console to create and configure a message queue, send a message, receive and delete that message, and then delete the queue.

[Get started with the tutorial](#)

Amazon EventBridge

- **Get started with Amazon EventBridge**

The basis of EventBridge is to create rules that route events to a target. In this guide, you create a basic rule.

[Explore the guide](#)

- **Amazon EventBridge get started tutorials**

These tutorials will help you explore the features of EventBridge and how to use them.

[Get started with the tutorials](#)

- **Integrating with other AWS services**

The following tutorials show you how to integrate EventBridge with other AWS services.

[Get started with the tutorials](#)

- **Build event-driven architectures**

Learn the basics of event-driven design, how to choose the right AWS service for the job, as well as how to optimize for both cost and performance.

[Get started with the tutorial](#)

- **Building event-driven applications with Amazon EventBridge**

Learn how to build event-driven applications by connecting multiple applications, including SaaS applications and AWS services, using the serverless event bus provided by Amazon EventBridge.

[Get started with the tutorial](#)

- **Kafka connector for Amazon EventBridge**

This connector allows you to convert records from one or more Kafka topics into events, and send those events to the event bus of your choice.

[Explore the guide](#)

- **Introducing cross-account targets for Amazon EventBridge Event Buses**

Learn how you can use Event Buses to send events directly to targets in other AWS accounts, such as Amazon SQS, Amazon SNS, and AWS Lambda.

[Read the blog post](#)

Amazon MQ

- **Create a connected message broker**

Learn how to set up an Amazon MQ message broker and connect a Java application without rewriting your code.

[Get started with the tutorial](#)

- **Migrate to a message broker in the cloud**

Amazon MQ makes it easy to migrate to a message broker in the cloud, such as Apache ActiveMQ and RabbitMQ.

[Read the guide](#)

- **Creating and connecting to a RabbitMQ broker**

Learn how you can use the AWS Management Console to create a RabbitMQ broker and attach your application to it.

[Get started with the tutorial](#)

- **RabbitMQ workshop**

This workshop is a collection of labs covering different aspects and patterns of messaging using RabbitMQ.

[Get started with the workshop](#)

- **Introducing quorum queues in Amazon MQ for RabbitMQ**

Discover how using quorum queues with RabbitMQ can provide higher availability and data security.

[Read the blog post](#)

- **Creating and connecting to an ActiveMQ broker**

Learn how you can use the AWS Management Console to create a basic broker.

[Get started with the tutorial](#)

- **ActiveMQ workshop**

Explore messaging concepts such as queues, topics and features of Amazon MQ such as failover, network of brokers.

[Get started with the workshop](#)

- **Deploy and publish to an Amazon MQ broker using AWS serverless**

We walk you through deploying a serverless backend and an Amazon MQ broker in one step using the AWS SAM.

[Read the blog post](#)

- **Measuring Amazon MQ throughput using Maven 2 benchmark and AWS CDK**

Learn how to evaluate the throughput for Amazon MQ using the ActiveMQ Classic Maven Performance test plugin.

[Read the blog post](#)

Amazon Kinesis Data Streams

- **Getting started with Amazon Kinesis Data Streams**

Learn fundamental Kinesis Data Streams data flow principles and the steps necessary to put and get data from an Kinesis data stream.

[Explore the guide](#)

- **Build highly available streams with Amazon Kinesis Data Streams**

We compare and contrast different strategies for creating a highly available Kinesis data stream in case of service interruptions, delays, or outages in the primary Region of operation.

[Read the blog post](#)

- **Example Tutorials for Amazon Kinesis Data Streams**

These tutorials are designed to further assist you in understanding Amazon Kinesis Data Streams concepts and functionality.

[Get started with the tutorials](#)

- **Using AWS Lambda with Amazon Kinesis**

Learn how to create a Lambda function to consume events from a Kinesis stream.

[Get started with the tutorial](#)

- **Real Time Streaming with Amazon Kinesis**

Explore a series of lab exercises which help users build a streaming analytics application on AWS.

[Get started with the tutorials](#)

Amazon MSK

- **Getting started using Amazon MSK**

This tutorial shows you an example of how you can create an MSK cluster, produce and consume data, and monitor the health of your cluster using metrics.

[Get started with the tutorial](#)

- **Getting started using MSK Serverless clusters**

This tutorial shows you an example of how you can create an MSK Serverless cluster, create a client machine that can access it, and use the client to create topics on the cluster and to write data to those topics.

[Get started with the tutorial](#)

- **Amazon MSK Labs**

These labs can either be run on personal or corporate AWS accounts or accounts provisioned by AWS Account teams for events that use workshop studio.

[Get started with the labs](#)

- **Introducing Express brokers for Amazon MSK to deliver high throughput and faster scaling for your Kafka cluster**

Learn how Express brokers can lower costs, provide higher resiliency, and lower operational overhead for your Kafka workloads.

[Read the blog post](#)

AWS Step Functions

- **Getting started with AWS Step Functions**

These tutorials walk you through creating a basic workflow for processing credit card applications.

[Get started with the tutorials](#)

- **Introduction to Step Functions**

This course introduces the key components of Step Functions to help you get started managing workflows within an application.

[Get started with the course](#)

- **Large-scale data processing with AWS Step Functions**

Learn how large-scale data processing applications can be built using Step Functions.

[Get started with the workshop](#)

- **Design Patterns for AWS Step Functions**

Learn how to implement design patterns in your Step Functions state machines and why to use each one.

[Get started with the course](#)

- **Schedule a Serverless Workflow with AWS Step Functions and Amazon EventBridge Scheduler**

We show you how to invoke a state machine using EventBridge Scheduler based on the schedule you define.

[Get started with the tutorial](#)

- **AWS Step Functions Workshop**

Learn how to use the primary features of AWS Step Functions through a series of interactive modules.

[Get started with the workshop](#)

- **Securely share AWS resources across VPC and account boundaries with PrivateLink, VPC Lattice, EventBridge, and Step Functions**

Learn how to share AWS resources, such as EC2 instances and Amazon EKS container services, across Amazon VPC and AWS account boundaries with AWS Step Functions.

[Read the blog post](#)

- **Simplifying developer experience with variables and JSONata in AWS Step Functions**

Explore using variables and JSONata, an open source query and transformation language, to simplify data sharing across states and reduce the complexity of data manipulation.

[Read the blog post](#)

Amazon MWAA

- **Get started with Amazon Managed Workflows for Apache Airflow**

This guide describes the prerequisites and the required AWS resources needed to get started with Amazon MWAA.

[Explore the guide](#)

- **Configuring the `aws-mwaa-local-runner` in a CD pipeline**

This tutorial guides you through the process of building a continuous delivery (CD) pipeline in GitHub using Amazon Managed Workflows for Apache Airflow's `aws-mwaa-local-runner` to test your Apache Airflow code locally.

[Get started with the tutorial](#)

- **Restricting an Amazon MWAA user's access to a subset of DAGs**

We show how you can restrict individual Amazon MWAA users to only view and interact with a specific DAG or a set of DAGs.

[Get started with the tutorial](#)

- **Amazon MWAA for Analytics Workshop**

Learn to build and orchestrate data and ML pipelines that include many of the above mentioned services, and with that you will gain familiarity and a better understanding of the hooks and operators available as part of Airflow to manage your pipelines/workflows on AWS.

[Get started with the workshop](#)

Explore

Once you have determined which approach best fits your workload for your environment, we recommend that you review these resources to help you begin implementing your approach. You can find service-specific resources in the previous section, and general event-driven architecture resources in the following section.

- **Architecture diagrams**

Explore reference architecture diagrams to help you create a highly available, secure, flexible, and cost effective architectures.

[Explore architecture diagrams](#)

- **Whitepapers**

Explore whitepapers to help you get started, and learn best practices around event driven architectures.

[Explore whitepapers](#)

- **Blogs**

Explore blogs to help you stay up to date on the latest technologies, and modernize your applications.

[Explore blogs](#)

Document history

The following table describes the important changes to this decision guide. For notifications about updates to this guide, you can subscribe to an RSS feed.

Change	Description	Date
Guide updated	Updated to reflect new features and numerous editorial updates throughout.	April 16, 2025
Guide updated	Updated links.	June 25, 2024
Initial publication	Guide first published.	March 31, 2023