



Guide for AWS large migrations

AWS Prescriptive Guidance



AWS Prescriptive Guidance: Guide for AWS large migrations

Copyright © 2024 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

Introduction	1
Guidance for large migrations	1
About the tools	2
Phases of a large migration	3
Assess phase	4
Mobilize phase	4
Migrate and modernize phase	4
About workstreams	5
About the initialize and implement stages	5
Stage 1: Initializing	7
Step 1: Get ready	8
Step 2: Define rules	9
Step 3: Create portfolio runbooks	9
Step 4: Create migration runbooks	10
Step 5: Evaluate current state	10
Stage 2: Implementing	11
About the migration strategies	14
Retire	14
Retain	15
Rehost	16
Relocate	16
Repurchase	17
Replatform	18
Refactor or re-architect	19
Resources	20
AWS large migrations	20
Additional references	20
Contributors	21
Document history	22
Glossary	23
#	23
A	24
B	27
C	28

D	31
E	35
F	37
G	38
H	39
I	40
L	42
M	43
O	47
P	49
Q	51
R	51
S	54
T	57
U	58
V	59
W	59
Z	60

Guide for AWS large migrations

Amazon Web Services (AWS)

February 2022 ([document history](#))

Note

Before reading this guide, we recommend reading [Strategy and best practices for AWS large migrations](#). The strategy discusses best practices for large migrations and provides use cases from customers across various industries. This guide describes a high-level, phased approach for implementing those best practices.

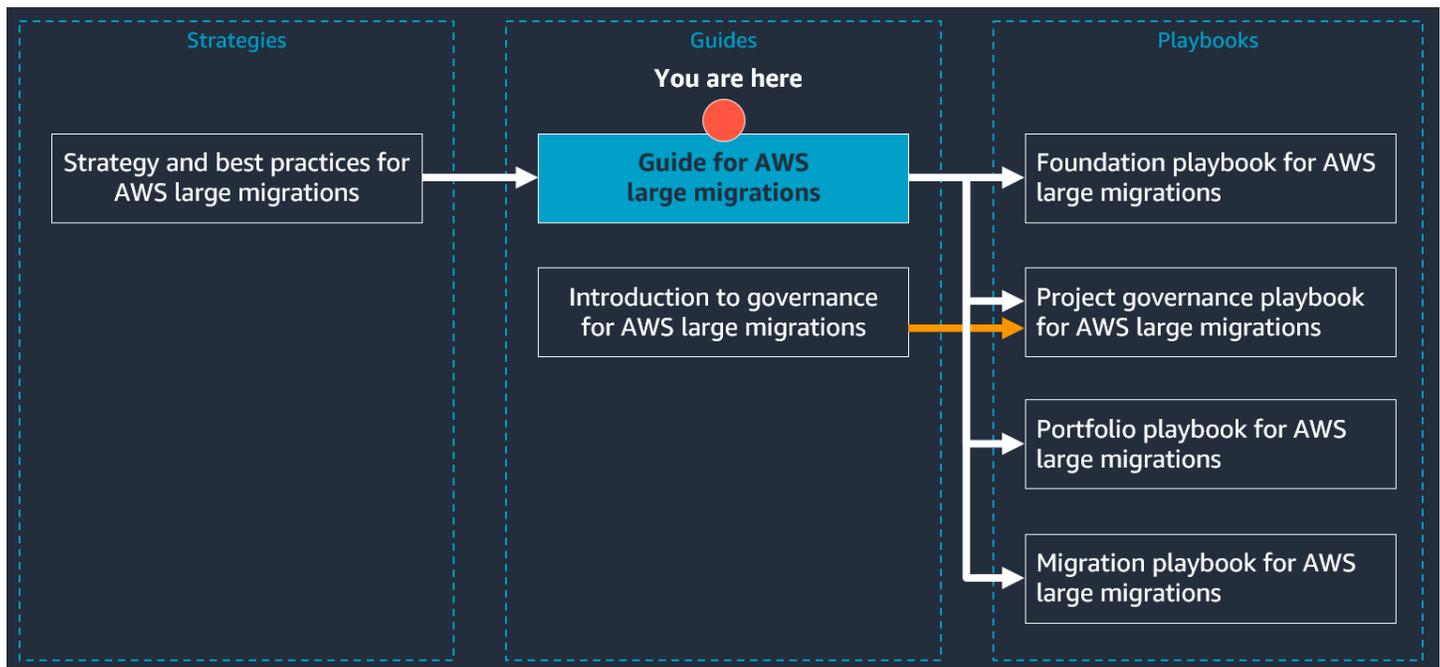
This guide provides an overview of the phased approach for large migrations, which consists of three phases: assess, mobilize, and migrate and modernize. The primary purpose of this guide is to describe the migrate part of the third phase, migrate and modernize, and its two stages, initialize and implement:

- In stage 1, *initialize*, you prepare your platform and people for a large migration. You also define the standard operating procedures (or *runbooks*) for the project. The runbooks and automations simplify and accelerate implementation of the large migration activities in stage 2.
- In stage 2, *implement*, you migrate servers at scale and closely manage and monitor the progress to implement continuous improvements.

Guidance for large migrations

Migrating 300 or more servers is considered a large migration. The people, process, and technology challenges of a large migration project are typically new to most enterprises. This document is part of an AWS Prescriptive Guidance series about large migrations to the AWS Cloud. This series is designed to help you apply the correct strategy and best practices from the outset, to streamline your journey to the cloud.

The following figure shows the other documents in this series. Review the strategy first, then the guides, and then proceed to the playbooks. To access the complete series, see [Large migrations to the AWS Cloud](#).



About the tools

This guide includes a [health-check matrix](#). You can use this tool to assess the health of your migration project throughout the migration. This tool helps you apply best practices and evaluate the efficiency and progress of your large migration project throughout its life cycle.

Additionally, each of the playbooks in this documentation series include templates and tools that can help you build each workflow.

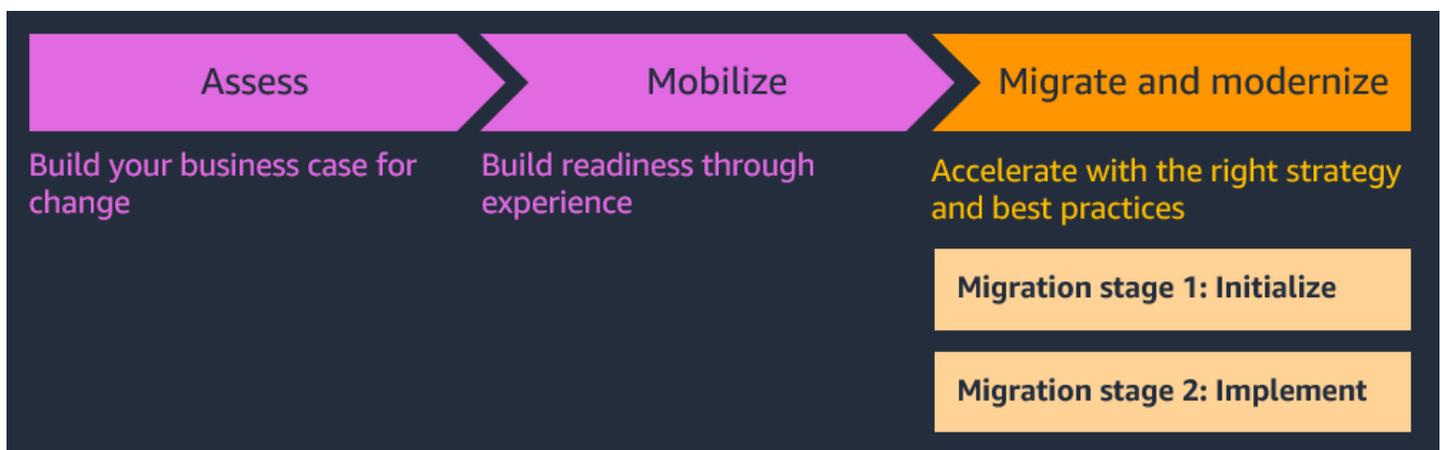
Phases of a large migration

AWS segments the large migration process into three sequential phases:

1. [Assess phase](#) – You build the business case for the migration.
2. [Mobilize phase](#) – You prepare the organization and mobilize the resources needed for the migration.
3. [Migrate and modernize phase of a large migration project](#) – You use your strategy, plan, and the best practices to migrate and modernize. In this phase, migration is divided into two stages, initialize and implement.

The assess and mobilize phases are the foundation of any large migration and prepare the organization for the migration. This guide includes a brief summary of the first two phases and provides references for additional support. The primary purpose of this guide is to describe the migrate part of the third phase, migrate and modernize, and its two stages, initialize and implement. We recommend that you first migrate to the AWS Cloud and then modernize the workload. Modernization is not included in this documentation set for large migrations. For more information about modernization, see [Phased approach to modernizing applications in the AWS Cloud](#) on the AWS Prescriptive Guidance website.

The following figure shows the three phases of a large migration: assess, mobilize, and migrate and modernize. Migration is divided into two stages, initialize and implement.



Even though this document set for large migration is focused on the migrate phase, we highly recommend you go through all three phases of a large migration in sequence. Completing the assess and mobilize phases builds a solid foundation to support the migration.

Assess phase

This is the start of your large-migration journey. In this phase, you build your business case for the large migration. You assess your organization's readiness for cloud transformation, take an initial look at your portfolio, and confirm that the key stakeholders in your organization are aligned.

For more information about the assess phase, see [Evaluating migration readiness](#) on the AWS Prescriptive Guidance website.

Mobilize phase

Now that you have built your business case for a cloud transformation, you use the readiness assessment from the previous phase and begin to fill in the gaps. You mobilize your organization and set up workstreams that prepare your organization for the large migration in the next phase. In this phase, you typically build your AWS landing zone, conduct a more thorough portfolio assessment, build your security and operating model, and prepare teams for change.

For more information about the mobilize phase, see [Mobilize your organization to accelerate large-scale migrations](#) on the AWS Prescriptive Guidance website.

Migrate and modernize phase of a large migration project

After the mobilize phase, your organization should have a solid foundation and be ready to start migration and modernization. Before proceeding, you evaluate the foundations, including the people and the platform, and make sure they are ready to support a large migration. The key for this step is to build standard operating procedures and automations that simplify and accelerate a repeatable pattern. This approach is known as a *migration factory*.

Just like a manufacturing factory, when you build a migration factory, you must initialize and calibrate it, define the standard operating procedures, measure the factory's performance, and continually improve the process and tools. In a large migration, initialization is a critical first stage. Running and improving the factory is the second stage, implementation, which migrates your servers at scale in batches known as waves. You organize your resources into workstreams, which are focused on accomplishing one aspect of the large migration.

This section consists of the following topics:

- [About workstreams](#)

- [About the initialize and implement stages](#)

About workstreams

Each workstream in your large migration project is dedicated to completing certain tasks, and those tasks change as you progress from stage 1, initialization, into stage 2, implementation. Although each workstream is independent, they work together to accomplish the same goal—migrate servers at scale.

The following are the four core workstreams, and you can create additional, supporting workstreams as needed to support your use case:

- **Foundation workstream** – This workstream is focused on preparing the people and platform for the large migration.
- **Project governance workstream** – This workstream manages the overall migration project, facilitates communication, and focuses on completing the project within budget and on time.
- **Portfolio workstream** – The teams in this workstream collect metadata to support the migration, prioritize applications, and perform wave planning.
- **Migration workstream** – Using the wave plan and collected metadata from the portfolio workstream, the teams in this workstream migrate and cutover the applications and servers.

For more information and instructions about how to establish and resource the workstreams in your large migration project, see the [Foundation playbook for AWS large migrations](#).

About the initialize and implement stages

Migrate and modernize is the third and final phase of your migration journey. Generally, you migrate first and then modernize. Migrate consists of two stages:

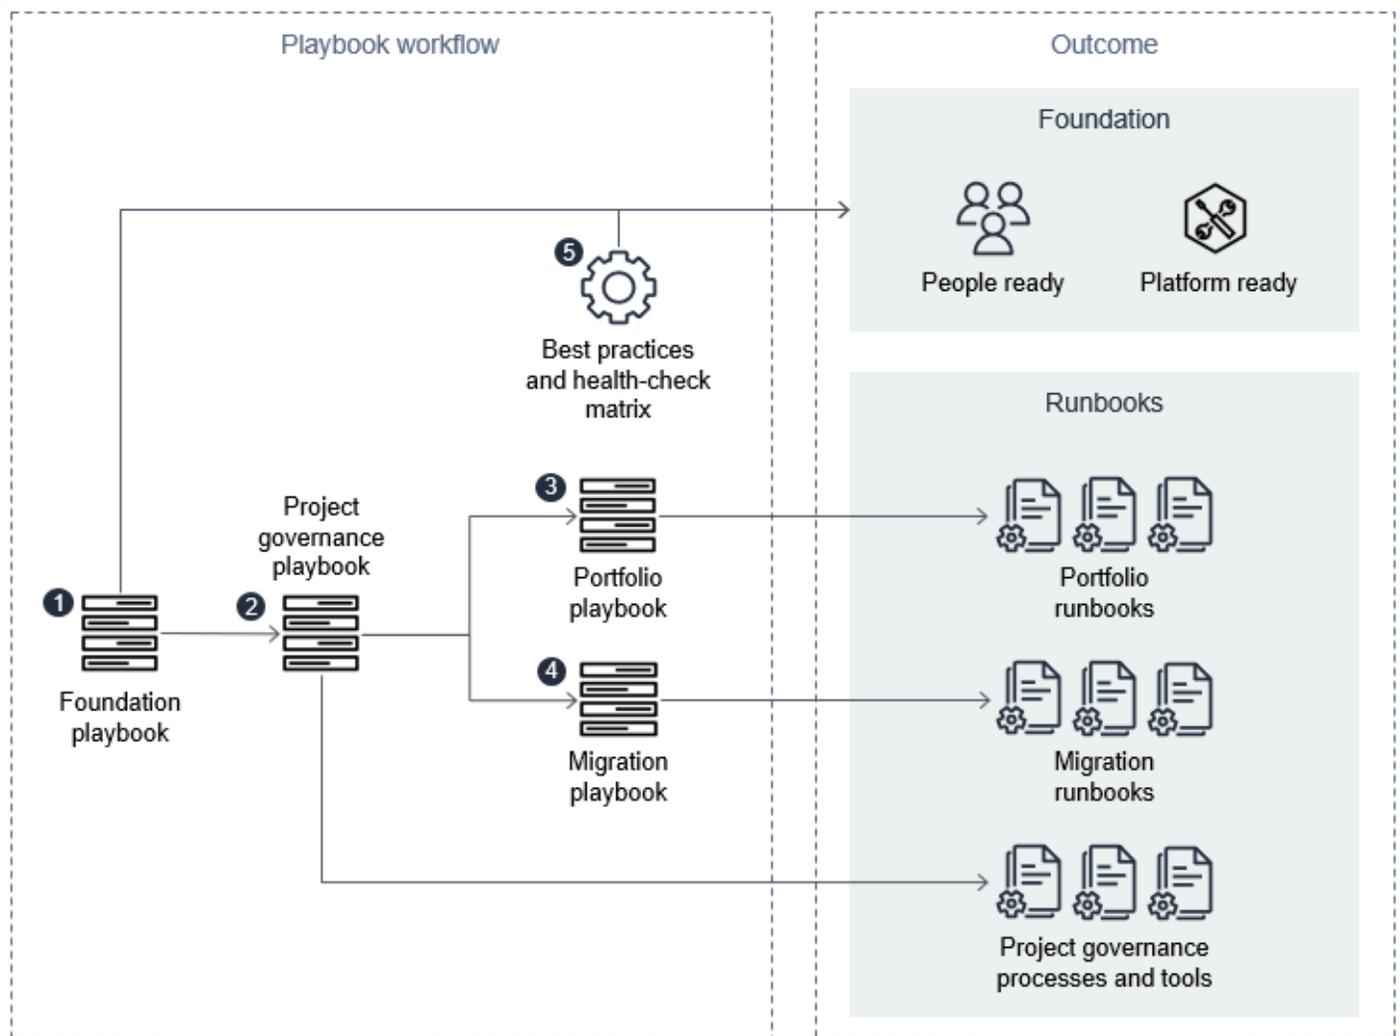
- **Stage 1: Initialize a large migration** – You prepare your platform and people for a large migration. For example, you review the landing zone design and verify that it has enough bandwidth to support the scale, and you design and implement a training plan. In this stage, based on your organization's policies and processes, you also define the standard operating procedures (or *runbooks*) for the large migration. The runbooks and automations simplify and accelerate implementation of the large migration in stage 2.
- **Stage 2: Implement a large migration** – In this stage, you migrate servers at scale by using the runbooks defined in stage 1. You manage the migration factory with project governance tools,

monitor the migration with a health-check matrix, and continuously improve the runbooks in order to increase the velocity of the migration.

Stage	Duration	Purpose
Stage 1: Initialize	1–3 months	<ul style="list-style-type: none">• Prepare your platform and people for a large migration.• Build your standard operating procedures (runbooks).
Stage 2: Implement	Varies by project scope and strategy	<ul style="list-style-type: none">• Use runbooks to implement the large migration.• Manage, monitor, and improve the migration.

Stage 1: Initializing a large migration

In the initialize stage, it is critical to define the standard operating procedures, or *runbooks*, for most of the migration tasks. Because there are many people, processes, and tools involved in a large migration, you can improve efficiency by defining, connecting, and automating procedures in a runbook. Runbooks provide clear guidance that everyone must follow. Implementing a large migration without runbooks increases the complexity of the migration and might cause the migration effort to lose momentum or fail. To help you get started, this guide reviews the workflow and playbooks, which contain useful templates for building your own runbooks.



The playbooks help you complete the first stage, *initialize*, and prepare for the second stage, *implement*:

1. [Foundation playbook for AWS large migrations](#) – This playbook helps you prepare your infrastructure and people.
2. [Project governance playbook for AWS large migrations](#) – This playbook helps you build project governance processes and tools that effectively manage a large migration project. It includes templates to help you get started.
3. [Portfolio playbook for AWS large migrations](#) – This playbook helps you build your own portfolio runbooks for wave planning and portfolio assessment. It includes templates to help you get started.
4. [Migration playbook for AWS large migrations](#) – This playbook helps you build a runbook for each migration pattern. It includes several templates for different patterns.
5. [Health-check matrix for AWS large migrations](#) – This helps you apply best practices and evaluate the efficiency and progress of your large migration throughout the life cycle of the project.

Step 1: Use the foundation playbook to get ready

In this playbook, you get your platform and people ready.

The following are examples of how you might prepare the people in your organization:

- Make sure your stakeholders are committed to the migration and ready to provide support.
- Define the workstreams in your large migration.
- Define your Cloud Enablement Engine, also known as *Cloud Center of Excellence*, team.
- Create a responsibility assignment model, in the form of a responsible, accountable, consulted, informed (RACI) matrix.
- Establish a training plan for your migration resources.

The following are examples of how you might prepare your organization's platform:

- Get the AWS landing zone ready to support a large migration.
- Evaluate the readiness of your on-premises infrastructure.

For more information, see the [Foundation playbook for AWS large migrations](#).

Step 2: Use the project governance playbook to define the rules

In this step, you define the rules, boundaries, and plans for running the large migration. This includes defining the following plans and processes:

- Communication plan
- Benefit-tracking office
- Escalation process
- Migration and cutover communication gates
- Project management processes and tools

For more information, see [Introduction to governance for AWS large migrations](#) and the [Project governance playbook for AWS large migrations](#).

Step 3: Use the portfolio playbook to create portfolio runbooks

Now, you get a better understanding of the scope and strategy. In this step, you follow the portfolio playbook to build detailed runbooks for portfolio assessment and wave planning. The goal is to make sure the migration factory doesn't run out of raw materials (servers) and has all the information needed to support the migration. Tasks include the following:

- Review the migration strategy from the mobilize phase (if you defined it in that phase). For more information, see [About the migration strategies](#).
- Define metadata requirements for each migration pattern.
- Define the metadata collection process and a metadata storage location that serves as a single source of truth.
- Define application prioritization rules.
- Define application deep dive rules.
- Define the wave planning process.

For more information, see the [Portfolio playbook for AWS large migrations](#).

Step 4: Use the migration playbook to build a migration runbook for each pattern

This step can run in parallel with the previous step. Using the migration playbook, you build your runbook for each of the migration patterns, such as rehosting to Amazon Elastic Compute Cloud (Amazon EC2) with AWS Cloud Migration Factory, or replatforming storage to Amazon Elastic File System (Amazon EFS) with AWS DataSync. You repeatedly test the runbooks until there are minimal to no errors. Tasks include the following:

- Understand the migration patterns.
- Get the runbook templates and add your on-premises process and tools.
- Test and improve the runbooks.
- Automate the runbooks as much as possible.

For more information, see the [Migration playbook for AWS large migrations](#).

Step 5: Use the health-check matrix to evaluate your current state

In this final step, you verify that your project team has completed the previous steps and identify any remaining gaps. Evaluating the health of the migration and improving the process are critical to the migration's success and helps maintain alignment with best practices.

Use the [health-check matrix](#) to evaluate the current state of your migration from a people, process, and technology perspective. This is not a one-off process. You should perform this evaluation on a regular basis.

Stage 2: Implementing a large migration

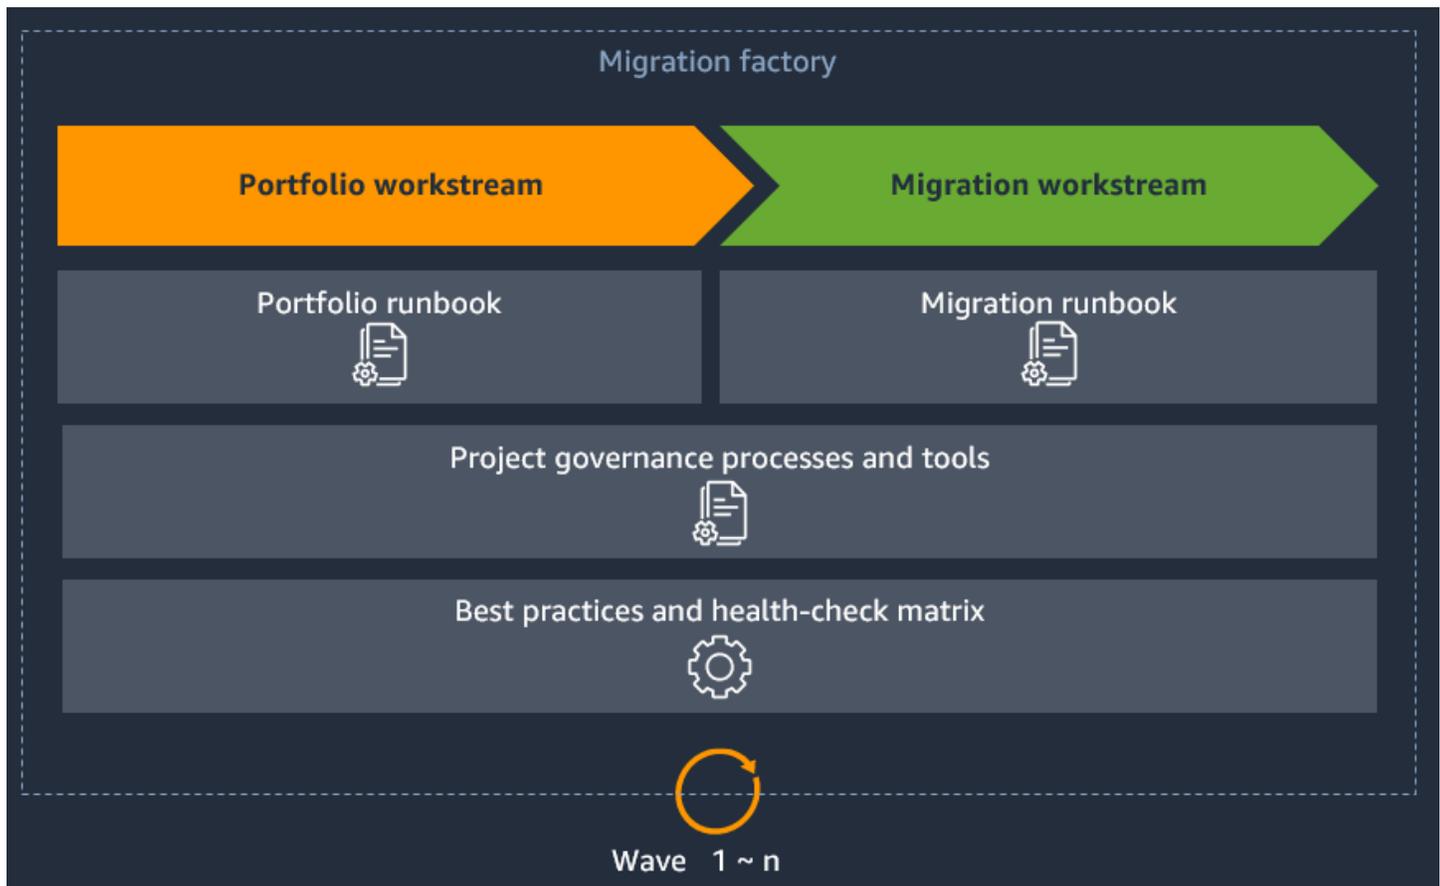
In stage 2 of a large migration, the goal is to migrate your servers at scale. For example, to migrate 1,000 servers in 6 months, you might start by migrating 5 servers per week and then gradually increase the velocity up to 50–100 servers per week.

Now you use the runbooks that you developed in stage 1 to migrate servers in waves. The first couple of waves are typically small because the migration and portfolio workstreams are adopting and adjusting the processes in their runbooks. Improving the runbooks is key to the success of a large migration. Runbooks are living documents. You must review, revise, and improve your runbooks after each cutover. As the runbooks evolve over time, the velocity should increase with every wave.

In stage 2, you use the following components to operate the migration factory:

- **Project governance rules** – You follow the project governance processes to manage waves, communication, timelines, and cutovers. These processes and tools make sure that everyone does the right thing at the right time and in the right order.
- **Portfolio runbooks** – You use the portfolio runbooks to prioritize applications, plan waves, and collect the required metadata that supports the migration. This metadata is the equivalent of raw materials in a manufacturing factory.
- **Migration runbooks** – You use the migration runbooks to migrate apps and servers, load the metadata to your migration tools, and complete the cutover process at the end of each wave. When following the migration runbooks, you adhere to the wave plan in the portfolio runbooks and use the metadata in the portfolio runbooks or from another single source of truth.
- **Large migration best practices and health-check matrix** – You use the health-check matrix to evaluate your current state frequently and regularly to make sure that everything is on track.

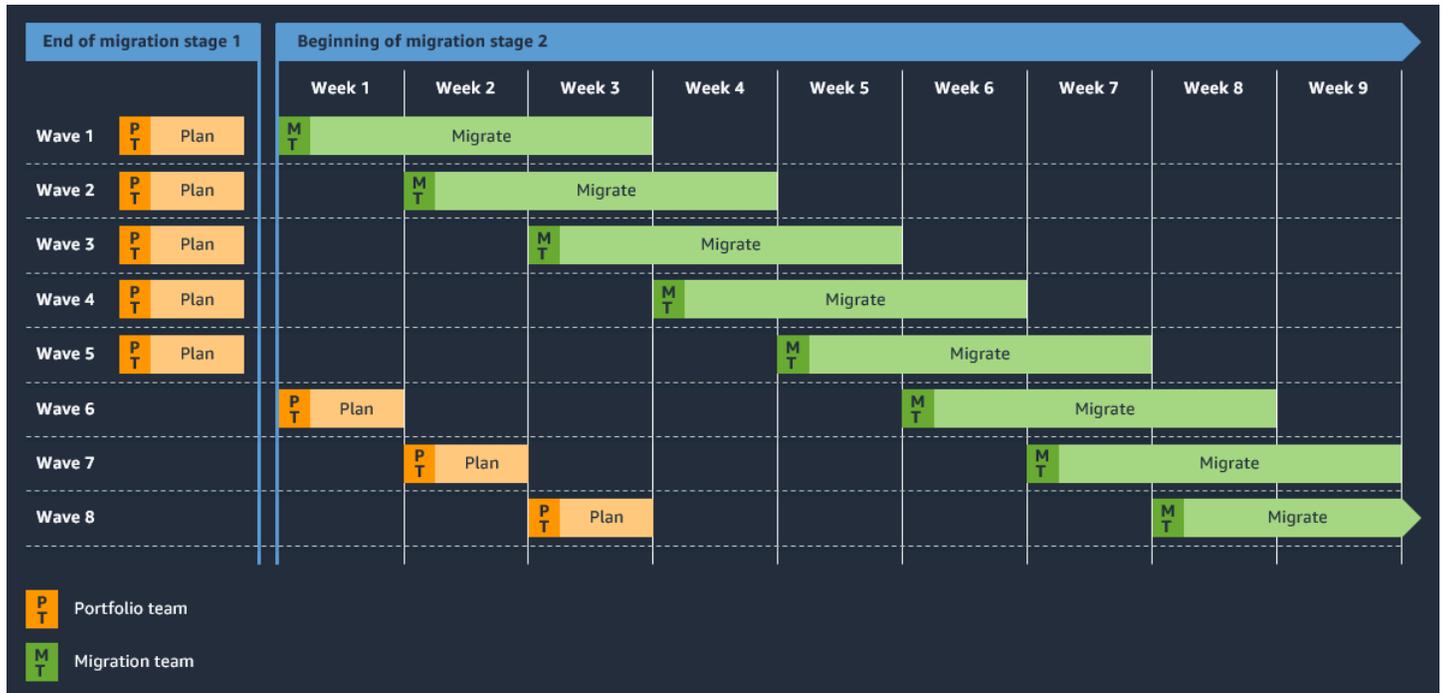
The following figure shows a typical migration factory for large migrations.



The runbooks are the key components of the migration factory, and they work together to form a data flow through two workstreams, portfolio and migration. For more information about these workstreams, see the [Foundation playbook for AWS large migrations](#). Rather than seeing a wave all the way through the migration factory, teams are usually dedicated to certain parts of the factory, and the waves flow through each workstream. The duration of each workstream varies based on your project timeline, scope, and resource availability. For example, the portfolio workstream might be 3 weeks, and the migration workstream might be 2–5 weeks. Prevent supply chain problems in your migration factory by ensuring that there are sufficient server waves lined up for migration. We recommend that the portfolio workstream is five waves ahead of the migration workstream.

The following figure shows a dynamic view of a typical migration factory. For each wave, the portfolio workstream runs 1–2 weeks, and the migration workstream typically runs 3–4 weeks. The portfolio workstream is five waves ahead of the migration workstream, so there is always a five-wave buffer between the portfolio and migration workstreams. At the end of migration stage 1, initializing, the portfolio workstream completes wave planning for a buffer of five waves. When the migration workstream starts migrating applications, this indicates that you have entered stage 2,

implementing. Both the portfolio and the migration workstreams continue to process waves, and the buffer prevents the migration workstream from running out of servers to migrate.



About the migration strategies

A *migration strategy* is the approach used to migrate a workload into the AWS Cloud. There are seven migration strategies for moving applications to the cloud, known as the 7 Rs:

- [Retire](#)
- [Retain](#)
- [Rehost](#)
- [Relocate](#)
- [Repurchase](#)
- [Replatform](#)
- [Refactor or re-architect](#)

Common strategies for large migrations include rehost, replatform, relocate, and retire. Refactor is not recommended for large migrations because it involves modernizing the application during the migration. This is the most complex of the migration strategies, and it can be complicated to manage for a large number of applications. Instead, we recommend rehosting, relocating, or replatforming the application and then modernizing the application after the migration is complete.

Selecting migration strategies is critical to a large migration. You might have selected migration strategies in the mobilize phase or during the initial portfolio assessment. This section reviews each migration strategy and their common use cases.

Retire

This is the migration strategy for the applications that you want to decommission or archive. Retiring the application means that you can shut down the servers within that application stack. The following are common use cases for the retire strategy:

- There is no business value in retaining the application or moving it to cloud.
- You want to eliminate the cost of maintaining and hosting the application.
- You want to reduce the security risks of operating an application that uses an operating system (OS) version or components that are no longer supported.

- You might want to retire applications based on their performance. For example, you might want to retire applications that have an average CPU and memory usage below 5 percent, known as *zombie applications*. You might also choose to retire some applications that have an average CPU and memory usage between 5 and 20 percent over a period of 90 days, known as *idle applications*. You can use the utilization and performance data from your discovery tool to identify zombie and idle applications.
- There has been no inbound connection to the application for the last 90 days.

For more information, see [Best practices for assessing applications to be retired during a migration to the AWS Cloud](#).

Retain

This is the migration strategy for applications that you want to keep in your source environment or applications that you are not ready to migrate. You might choose to migrate these applications in the future.

The following are common use cases for the retain strategy:

- **Security and compliance** – You might want to retain applications in order to remain in compliance with data residency requirements.
- **High risk** – You might decide to retain an application because it requires a detailed assessment and plan prior to migration.
- **Dependencies** – You might decide to retain an application if you need to migrate one or more other applications first.
- **Applications that are recently upgraded** – You might want to postpone migrating the application until the next technical refresh because you recently invested in upgrading your current system.
- **No business value to migrate** – There is no business value for migrating some applications to the cloud, such as those with only a few internal users.
- **Plans to migrate to software as a service (SaaS)** – You might choose retain an application until the SaaS version is released by the vendor. This is a common strategy for vendor-based applications.
- **Unresolved physical dependencies** – You might choose to retain an application that is dependent on specialized hardware that does not have a cloud equivalent, such as machines in a manufacturing plant.

- **Mainframe or mid-range applications and non-x86 Unix applications** – These applications require careful assessment and planning before migrating them to the cloud. Examples of mid-range applications include IBM AS/400 and Oracle Solaris.
- **Performance** – You might want to retain applications based on their performance. For example, you might want to keep zombie or idle applications in your source environment.

Rehost

This strategy is also known as *lift and shift*. Using this strategy, you move your applications from your source environment to the AWS Cloud without making any changes to the application. For example, you migrate your application stack from on-premises to the AWS Cloud.

With rehost, you can migrate a large number of machines from multiple source platforms (physical, virtual, or another cloud) to the AWS Cloud without worrying about compatibility, performance disruption, long cutover windows, or long-distance data replications.

Your application continues to serve users while the workloads are being migrated, which minimizes disruption and downtime. The downtime depends on your cutover strategy.

This strategy helps you to scale your applications without implementing any cloud optimizations that could save you time or money. Applications are easier to optimize or re-architect when they are already running in cloud because it is easier to integrate to AWS services and manage your workloads.

You can automate rehosting by using the following services:

- [AWS Application Migration Service](#)
- [AWS Cloud Migration Factory Solution](#)
- [VM Import/Export](#)

For a list of migration patterns for the rehost migration strategy, see [Rehost](#) on the AWS Prescriptive Guidance website.

Relocate

Using this strategy, you can transfer a large number of servers, comprising one or more applications, at a given time from on-premises platform to a cloud version of the platform. You

can also use the relocate strategy to move instances or objects to a different virtual private cloud (VPC), AWS Region, or AWS account. For example, you can use this strategy to transfer servers in bulk from VMware software-defined data center (SSDC) to VMware Cloud on AWS, or you can transfer an Amazon Relational Database Service (Amazon RDS) DB instance to another VPC or AWS account.

The relocate strategy doesn't require that you purchase new hardware, rewrite applications, or modify your existing operation. During relocation, the application continues to serve users, which minimizes disruption and downtime. Relocate is the quickest way to migrate and operate your workload in the cloud because it does not impact the overall architecture of your application.

For a list of migration patterns for the relocate migration strategy, see [Relocate](#) on the AWS Prescriptive Guidance website.

Repurchase

This strategy is also known as *drop and shop*. You replace your application with a different version or product. The new application should provide more business value than the existing, on-premises application, including features such as accessibility from anywhere, no infrastructure to maintain, and pay-as-you-go pricing models. Repurchasing the application typically reduces costs associated with maintenance, infrastructure, and licensing.

The following are common use cases for the repurchase migration strategy:

- **Moving from a traditional license to SaaS** – This removes the burden of managing and maintaining the infrastructure and helps reduce licensing issues.
- **Version upgrades or third-party equivalents** – By replacing your existing on-premises application with the vendor's latest version or third-party equivalent in the cloud, you can leverage new features, integrate with cloud services, and scale the application more easily.
- **Replacing a custom application** – You can avoid recoding and re-architecting a custom application by repurchasing a vendor-based SaaS or cloud-based application.

Before purchasing, you need to assess the application according to your business requirements, especially security and compliance.

After you purchase the new application, the following are the next steps:

- Training your team and users with the new system

- Migrating your data to the newly purchased application
- Integrating the application into your authentication services, such as Microsoft Active Directory, to centralize authentication
- Configuring networking to help secure communication between the purchased application, your users, and your infrastructure

Typically, the application vendor helps you with these activities for a smooth transition.

Replatform

This strategy is also known as *lift, tinker, and shift* or *lift and reshape*. Using this migration strategy, you move the application to the cloud, and you introduce some level of optimization in order to operate the application efficiently, to reduce costs, or to take advantage of cloud capabilities. For example, you might replatform a Microsoft SQL Server database to Amazon RDS for SQL Server.

Using this strategy, you might make a few or many changes to the application, depending on your business goals and your target platform.

The following are common use cases for the replatform migration strategy:

- You want to save time and reduce cost by moving to a fully managed service or serverless service in the AWS Cloud.
- You want to improve your security and compliance stance by upgrading your operating system to the latest version. By using [End-of-Support Migration Program \(EMP\) for Windows Server](#), you can migrate your legacy Windows Server applications to the latest, supported versions of Windows Server on AWS, without any code changes. You can use this [decision tree](#) in the *AWS EMP for Windows Server User Guide* to help you determine your EMP workloads.
- You can reduce costs by using [AWS Graviton Processors](#), custom-built processors developed by AWS.
- You can reduce costs by moving from a Microsoft Windows operating system to a Linux operating system. You can port your .NET Framework applications to .NET Core, which can run on a Linux operating system. [Porting Assistant for .NET](#) is an analysis tool that helps you port your applications to Linux.
- You can improve performance by migrating virtual machines into containers, without making any code changes. You can modernize your .NET and Java applications into containerized applications by using the [AWS App2Container migration tool](#).

The replatform strategy keeps your legacy application running without compromising security and compliance.

Replatform reduces cost and improves performance by migrating to a managed or serverless service, moving virtual machines to container, and avoiding licensing expenses.

For a list of migration patterns for the replatform migration strategy, see [Replatform](#) on the AWS Prescriptive Guidance website.

Refactor or re-architect

Using this strategy, you move an application to the cloud and modify its architecture by taking full advantage of cloud-native features to improve agility, performance, and scalability. This is driven by strong business demand to scale, accelerate product and feature releases, and to reduce costs.

The following are common use cases for the refactor migration strategy:

- The legacy mainframe application can no longer address the demand of the business due to its limitations or is expensive to maintain.
- You have a monolith application that is already hindering efforts to deliver product quickly or address customer needs and demands.
- You have a legacy application that nobody knows how to maintain, or the source code is unavailable.
- The application is difficult to test, or test coverage is very low. This affects the quality and delivery of new application features and fixes. By redesigning the application for the cloud, you can increase the test coverage and integrate automated testing tools.
- For security and compliance reasons, when moving a database to the cloud, you might need to extract some tables (such as customer information, patient, or patient diagnosis tables) and retain those tables on premises. In this situation, you need to refactor your database in order to separate the tables that will be migrated from those that will be retained on premises.

For a list of migration patterns for the refactor migration strategy, see [Re-architect](#) on the AWS Prescriptive Guidance website.

Resources

AWS large migrations

To access the complete AWS Prescriptive Guidance series for large migrations, see [Large migrations to the AWS Cloud](#).

Additional references

- [AWS Cloud Migration Factory Solution](#)
- [When to Choose AWS Application Migration Service](#)
- [Best practices for assessing applications to be retired during a migration to the AWS Cloud](#)
- [AWS Prescriptive Guidance migration patterns](#)

Contributors

The following individuals contributed to this document:

- Wally Lu, Principal Consultant
- Damien Renner, Senior Migration Specialist

Document history

The following table describes significant changes to this guide. If you want to be notified about future updates, you can subscribe to an [RSS feed](#).

Change	Description	Date
Updated name of AWS solution	We updated the name of the referenced AWS solution from <i>CloudEndure Migration Factory</i> to <i>Cloud Migration Factory</i> .	May 2, 2022
Initial publication	—	February 28, 2022

AWS Prescriptive Guidance glossary

The following are commonly used terms in strategies, guides, and patterns provided by AWS Prescriptive Guidance. To suggest entries, please use the **Provide feedback** link at the end of the glossary.

Numbers

7 Rs

Seven common migration strategies for moving applications to the cloud. These strategies build upon the 5 Rs that Gartner identified in 2011 and consist of the following:

- **Refactor/re-architect** – Move an application and modify its architecture by taking full advantage of cloud-native features to improve agility, performance, and scalability. This typically involves porting the operating system and database. Example: Migrate your on-premises Oracle database to the Amazon Aurora PostgreSQL-Compatible Edition.
- **Replatform (lift and reshape)** – Move an application to the cloud, and introduce some level of optimization to take advantage of cloud capabilities. Example: Migrate your on-premises Oracle database to Amazon Relational Database Service (Amazon RDS) for Oracle in the AWS Cloud.
- **Repurchase (drop and shop)** – Switch to a different product, typically by moving from a traditional license to a SaaS model. Example: Migrate your customer relationship management (CRM) system to Salesforce.com.
- **Rehost (lift and shift)** – Move an application to the cloud without making any changes to take advantage of cloud capabilities. Example: Migrate your on-premises Oracle database to Oracle on an EC2 instance in the AWS Cloud.
- **Relocate (hypervisor-level lift and shift)** – Move infrastructure to the cloud without purchasing new hardware, rewriting applications, or modifying your existing operations. This migration scenario is specific to VMware Cloud on AWS, which supports virtual machine (VM) compatibility and workload portability between your on-premises environment and AWS. You can use the VMware Cloud Foundation technologies from your on-premises data centers when you migrate your infrastructure to VMware Cloud on AWS. Example: Relocate the hypervisor hosting your Oracle database to VMware Cloud on AWS.
- **Retain (revisit)** – Keep applications in your source environment. These might include applications that require major refactoring, and you want to postpone that work until a later

time, and legacy applications that you want to retain, because there's no business justification for migrating them.

- **Retire** – Decommission or remove applications that are no longer needed in your source environment.

A

ABAC

See [attribute-based access control](#).

abstracted services

See [managed services](#).

ACID

See [atomicity, consistency, isolation, durability](#).

active-active migration

A database migration method in which the source and target databases are kept in sync (by using a bidirectional replication tool or dual write operations), and both databases handle transactions from connecting applications during migration. This method supports migration in small, controlled batches instead of requiring a one-time cutover. It's more flexible but requires more work than [active-passive migration](#).

active-passive migration

A database migration method in which in which the source and target databases are kept in sync, but only the source database handles transactions from connecting applications while data is replicated to the target database. The target database doesn't accept any transactions during migration.

aggregate function

A SQL function that operates on a group of rows and calculates a single return value for the group. Examples of aggregate functions include SUM and MAX.

AI

See [artificial intelligence](#).

AIOps

See [artificial intelligence operations](#).

anonymization

The process of permanently deleting personal information in a dataset. Anonymization can help protect personal privacy. Anonymized data is no longer considered to be personal data.

anti-pattern

A frequently used solution for a recurring issue where the solution is counter-productive, ineffective, or less effective than an alternative.

application control

A security approach that allows the use of only approved applications in order to help protect a system from malware.

application portfolio

A collection of detailed information about each application used by an organization, including the cost to build and maintain the application, and its business value. This information is key to [the portfolio discovery and analysis process](#) and helps identify and prioritize the applications to be migrated, modernized, and optimized.

artificial intelligence (AI)

The field of computer science that is dedicated to using computing technologies to perform cognitive functions that are typically associated with humans, such as learning, solving problems, and recognizing patterns. For more information, see [What is Artificial Intelligence?](#)

artificial intelligence operations (AIOps)

The process of using machine learning techniques to solve operational problems, reduce operational incidents and human intervention, and increase service quality. For more information about how AIOps is used in the AWS migration strategy, see the [operations integration guide](#).

asymmetric encryption

An encryption algorithm that uses a pair of keys, a public key for encryption and a private key for decryption. You can share the public key because it isn't used for decryption, but access to the private key should be highly restricted.

atomicity, consistency, isolation, durability (ACID)

A set of software properties that guarantee the data validity and operational reliability of a database, even in the case of errors, power failures, or other problems.

attribute-based access control (ABAC)

The practice of creating fine-grained permissions based on user attributes, such as department, job role, and team name. For more information, see [ABAC for AWS](#) in the AWS Identity and Access Management (IAM) documentation.

authoritative data source

A location where you store the primary version of data, which is considered to be the most reliable source of information. You can copy data from the authoritative data source to other locations for the purposes of processing or modifying the data, such as anonymizing, redacting, or pseudonymizing it.

Availability Zone

A distinct location within an AWS Region that is insulated from failures in other Availability Zones and provides inexpensive, low-latency network connectivity to other Availability Zones in the same Region.

AWS Cloud Adoption Framework (AWS CAF)

A framework of guidelines and best practices from AWS to help organizations develop an efficient and effective plan to move successfully to the cloud. AWS CAF organizes guidance into six focus areas called perspectives: business, people, governance, platform, security, and operations. The business, people, and governance perspectives focus on business skills and processes; the platform, security, and operations perspectives focus on technical skills and processes. For example, the people perspective targets stakeholders who handle human resources (HR), staffing functions, and people management. For this perspective, AWS CAF provides guidance for people development, training, and communications to help ready the organization for successful cloud adoption. For more information, see the [AWS CAF website](#) and the [AWS CAF whitepaper](#).

AWS Workload Qualification Framework (AWS WQF)

A tool that evaluates database migration workloads, recommends migration strategies, and provides work estimates. AWS WQF is included with AWS Schema Conversion Tool (AWS SCT). It analyzes database schemas and code objects, application code, dependencies, and performance characteristics, and provides assessment reports.

B

BCP

See [business continuity planning](#).

behavior graph

A unified, interactive view of resource behavior and interactions over time. You can use a behavior graph with Amazon Detective to examine failed logon attempts, suspicious API calls, and similar actions. For more information, see [Data in a behavior graph](#) in the Detective documentation.

big-endian system

A system that stores the most significant byte first. See also [endianness](#).

binary classification

A process that predicts a binary outcome (one of two possible classes). For example, your ML model might need to predict problems such as "Is this email spam or not spam?" or "Is this product a book or a car?"

bloom filter

A probabilistic, memory-efficient data structure that is used to test whether an element is a member of a set.

branch

A contained area of a code repository. The first branch created in a repository is the *main branch*. You can create a new branch from an existing branch, and you can then develop features or fix bugs in the new branch. A branch you create to build a feature is commonly referred to as a *feature branch*. When the feature is ready for release, you merge the feature branch back into the main branch. For more information, see [About branches](#) (GitHub documentation).

break-glass access

In exceptional circumstances and through an approved process, a quick means for a user to gain access to an AWS account that they don't typically have permissions to access. For more information, see the [Implement break-glass procedures](#) indicator in the AWS Well-Architected guidance.

brownfield strategy

The existing infrastructure in your environment. When adopting a brownfield strategy for a system architecture, you design the architecture around the constraints of the current systems and infrastructure. If you are expanding the existing infrastructure, you might blend brownfield and [greenfield](#) strategies.

buffer cache

The memory area where the most frequently accessed data is stored.

business capability

What a business does to generate value (for example, sales, customer service, or marketing). Microservices architectures and development decisions can be driven by business capabilities. For more information, see the [Organized around business capabilities](#) section of the [Running containerized microservices on AWS](#) whitepaper.

business continuity planning (BCP)

A plan that addresses the potential impact of a disruptive event, such as a large-scale migration, on operations and enables a business to resume operations quickly.

C

CAF

See [AWS Cloud Adoption Framework](#).

CCoE

See [Cloud Center of Excellence](#).

CDC

See [change data capture](#).

change data capture (CDC)

The process of tracking changes to a data source, such as a database table, and recording metadata about the change. You can use CDC for various purposes, such as auditing or replicating changes in a target system to maintain synchronization.

chaos engineering

Intentionally introducing failures or disruptive events to test a system's resilience. You can use [AWS Fault Injection Service \(AWS FIS\)](#) to perform experiments that stress your AWS workloads and evaluate their response.

CI/CD

See [continuous integration and continuous delivery](#).

classification

A categorization process that helps generate predictions. ML models for classification problems predict a discrete value. Discrete values are always distinct from one another. For example, a model might need to evaluate whether or not there is a car in an image.

client-side encryption

Encryption of data locally, before the target AWS service receives it.

Cloud Center of Excellence (CCoE)

A multi-disciplinary team that drives cloud adoption efforts across an organization, including developing cloud best practices, mobilizing resources, establishing migration timelines, and leading the organization through large-scale transformations. For more information, see the [CCoE posts](#) on the AWS Cloud Enterprise Strategy Blog.

cloud computing

The cloud technology that is typically used for remote data storage and IoT device management. Cloud computing is commonly connected to [edge computing](#) technology.

cloud operating model

In an IT organization, the operating model that is used to build, mature, and optimize one or more cloud environments. For more information, see [Building your Cloud Operating Model](#).

cloud stages of adoption

The four phases that organizations typically go through when they migrate to the AWS Cloud:

- Project – Running a few cloud-related projects for proof of concept and learning purposes
- Foundation – Making foundational investments to scale your cloud adoption (e.g., creating a landing zone, defining a CCoE, establishing an operations model)
- Migration – Migrating individual applications
- Re-invention – Optimizing products and services, and innovating in the cloud

These stages were defined by Stephen Orban in the blog post [The Journey Toward Cloud-First & the Stages of Adoption](#) on the AWS Cloud Enterprise Strategy blog. For information about how they relate to the AWS migration strategy, see the [migration readiness guide](#).

CMDB

See [configuration management database](#).

code repository

A location where source code and other assets, such as documentation, samples, and scripts, are stored and updated through version control processes. Common cloud repositories include GitHub or AWS CodeCommit. Each version of the code is called a *branch*. In a microservice structure, each repository is devoted to a single piece of functionality. A single CI/CD pipeline can use multiple repositories.

cold cache

A buffer cache that is empty, not well populated, or contains stale or irrelevant data. This affects performance because the database instance must read from the main memory or disk, which is slower than reading from the buffer cache.

cold data

Data that is rarely accessed and is typically historical. When querying this kind of data, slow queries are typically acceptable. Moving this data to lower-performing and less expensive storage tiers or classes can reduce costs.

computer vision

A field of AI used by machines to identify people, places, and things in images with accuracy at or above human levels. Often built with deep learning models, it automates extraction, analysis, classification, and understanding of useful information from a single image or a sequence of images.

configuration management database (CMDB)

A repository that stores and manages information about a database and its IT environment, including both hardware and software components and their configurations. You typically use data from a CMDB in the portfolio discovery and analysis stage of migration.

conformance pack

A collection of AWS Config rules and remediation actions that you can assemble to customize your compliance and security checks. You can deploy a conformance pack as a single entity in

an AWS account and Region, or across an organization, by using a YAML template. For more information, see [Conformance packs](#) in the AWS Config documentation.

continuous integration and continuous delivery (CI/CD)

The process of automating the source, build, test, staging, and production stages of the software release process. CI/CD is commonly described as a pipeline. CI/CD can help you automate processes, improve productivity, improve code quality, and deliver faster. For more information, see [Benefits of continuous delivery](#). CD can also stand for *continuous deployment*. For more information, see [Continuous Delivery vs. Continuous Deployment](#).

D

data at rest

Data that is stationary in your network, such as data that is in storage.

data classification

A process for identifying and categorizing the data in your network based on its criticality and sensitivity. It is a critical component of any cybersecurity risk management strategy because it helps you determine the appropriate protection and retention controls for the data. Data classification is a component of the security pillar in the AWS Well-Architected Framework. For more information, see [Data classification](#).

data drift

A meaningful variation between the production data and the data that was used to train an ML model, or a meaningful change in the input data over time. Data drift can reduce the overall quality, accuracy, and fairness in ML model predictions.

data in transit

Data that is actively moving through your network, such as between network resources.

data minimization

The principle of collecting and processing only the data that is strictly necessary. Practicing data minimization in the AWS Cloud can reduce privacy risks, costs, and your analytics carbon footprint.

data perimeter

A set of preventive guardrails in your AWS environment that help make sure that only trusted identities are accessing trusted resources from expected networks. For more information, see [Building a data perimeter on AWS](#).

data preprocessing

To transform raw data into a format that is easily parsed by your ML model. Preprocessing data can mean removing certain columns or rows and addressing missing, inconsistent, or duplicate values.

data provenance

The process of tracking the origin and history of data throughout its lifecycle, such as how the data was generated, transmitted, and stored.

data subject

An individual whose data is being collected and processed.

data warehouse

A data management system that supports business intelligence, such as analytics. Data warehouses commonly contain large amounts of historical data, and they are typically used for queries and analysis.

database definition language (DDL)

Statements or commands for creating or modifying the structure of tables and objects in a database.

database manipulation language (DML)

Statements or commands for modifying (inserting, updating, and deleting) information in a database.

DDL

See [database definition language](#).

deep ensemble

To combine multiple deep learning models for prediction. You can use deep ensembles to obtain a more accurate prediction or for estimating uncertainty in predictions.

deep learning

An ML subfield that uses multiple layers of artificial neural networks to identify mapping between input data and target variables of interest.

defense-in-depth

An information security approach in which a series of security mechanisms and controls are thoughtfully layered throughout a computer network to protect the confidentiality, integrity, and availability of the network and the data within. When you adopt this strategy on AWS, you add multiple controls at different layers of the AWS Organizations structure to help secure resources. For example, a defense-in-depth approach might combine multi-factor authentication, network segmentation, and encryption.

delegated administrator

In AWS Organizations, a compatible service can register an AWS member account to administer the organization's accounts and manage permissions for that service. This account is called the *delegated administrator* for that service. For more information and a list of compatible services, see [Services that work with AWS Organizations](#) in the AWS Organizations documentation.

deployment

The process of making an application, new features, or code fixes available in the target environment. Deployment involves implementing changes in a code base and then building and running that code base in the application's environments.

development environment

See [environment](#).

detective control

A security control that is designed to detect, log, and alert after an event has occurred. These controls are a second line of defense, alerting you to security events that bypassed the preventative controls in place. For more information, see [Detective controls](#) in *Implementing security controls on AWS*.

development value stream mapping (DVSM)

A process used to identify and prioritize constraints that adversely affect speed and quality in a software development lifecycle. DVSM extends the value stream mapping process originally designed for lean manufacturing practices. It focuses on the steps and teams required to create and move value through the software development process.

digital twin

A virtual representation of a real-world system, such as a building, factory, industrial equipment, or production line. Digital twins support predictive maintenance, remote monitoring, and production optimization.

dimension table

In a [star schema](#), a smaller table that contains data attributes about quantitative data in a fact table. Dimension table attributes are typically text fields or discrete numbers that behave like text. These attributes are commonly used for query constraining, filtering, and result set labeling.

disaster

An event that prevents a workload or system from fulfilling its business objectives in its primary deployed location. These events can be natural disasters, technical failures, or the result of human actions, such as unintentional misconfiguration or a malware attack.

disaster recovery (DR)

The strategy and process you use to minimize downtime and data loss caused by a [disaster](#). For more information, see [Disaster Recovery of Workloads on AWS: Recovery in the Cloud](#) in the AWS Well-Architected Framework.

DML

See [database manipulation language](#).

domain-driven design

An approach to developing a complex software system by connecting its components to evolving domains, or core business goals, that each component serves. This concept was introduced by Eric Evans in his book, *Domain-Driven Design: Tackling Complexity in the Heart of Software* (Boston: Addison-Wesley Professional, 2003). For information about how you can use domain-driven design with the strangler fig pattern, see [Modernizing legacy Microsoft ASP.NET \(ASMX\) web services incrementally by using containers and Amazon API Gateway](#).

DR

See [disaster recovery](#).

drift detection

Tracking deviations from a baselined configuration. For example, you can use AWS CloudFormation to [detect drift in system resources](#), or you can use AWS Control Tower to [detect changes in your landing zone](#) that might affect compliance with governance requirements.

DVSM

See [development value stream mapping](#).

E

EDA

See [exploratory data analysis](#).

edge computing

The technology that increases the computing power for smart devices at the edges of an IoT network. When compared with [cloud computing](#), edge computing can reduce communication latency and improve response time.

encryption

A computing process that transforms plaintext data, which is human-readable, into ciphertext.

encryption key

A cryptographic string of randomized bits that is generated by an encryption algorithm. Keys can vary in length, and each key is designed to be unpredictable and unique.

endianness

The order in which bytes are stored in computer memory. Big-endian systems store the most significant byte first. Little-endian systems store the least significant byte first.

endpoint

See [service endpoint](#).

endpoint service

A service that you can host in a virtual private cloud (VPC) to share with other users. You can create an endpoint service with AWS PrivateLink and grant permissions to other AWS accounts or to AWS Identity and Access Management (IAM) principals. These accounts or principals

can connect to your endpoint service privately by creating interface VPC endpoints. For more information, see [Create an endpoint service](#) in the Amazon Virtual Private Cloud (Amazon VPC) documentation.

envelope encryption

The process of encrypting an encryption key with another encryption key. For more information, see [Envelope encryption](#) in the AWS Key Management Service (AWS KMS) documentation.

environment

An instance of a running application. The following are common types of environments in cloud computing:

- development environment – An instance of a running application that is available only to the core team responsible for maintaining the application. Development environments are used to test changes before promoting them to upper environments. This type of environment is sometimes referred to as a *test environment*.
- lower environments – All development environments for an application, such as those used for initial builds and tests.
- production environment – An instance of a running application that end users can access. In a CI/CD pipeline, the production environment is the last deployment environment.
- upper environments – All environments that can be accessed by users other than the core development team. This can include a production environment, preproduction environments, and environments for user acceptance testing.

epic

In agile methodologies, functional categories that help organize and prioritize your work. Epics provide a high-level description of requirements and implementation tasks. For example, AWS CAF security epics include identity and access management, detective controls, infrastructure security, data protection, and incident response. For more information about epics in the AWS migration strategy, see the [program implementation guide](#).

exploratory data analysis (EDA)

The process of analyzing a dataset to understand its main characteristics. You collect or aggregate data and then perform initial investigations to find patterns, detect anomalies, and check assumptions. EDA is performed by calculating summary statistics and creating data visualizations.

F

fact table

The central table in a [star schema](#). It stores quantitative data about business operations. Typically, a fact table contains two types of columns: those that contain measures and those that contain a foreign key to a dimension table.

fail fast

A philosophy that uses frequent and incremental testing to reduce the development lifecycle. It is a critical part of an agile approach.

fault isolation boundary

In the AWS Cloud, a boundary such as an Availability Zone, AWS Region, control plane, or data plane that limits the effect of a failure and helps improve the resilience of workloads. For more information, see [AWS Fault Isolation Boundaries](#).

feature branch

See [branch](#).

features

The input data that you use to make a prediction. For example, in a manufacturing context, features could be images that are periodically captured from the manufacturing line.

feature importance

How significant a feature is for a model's predictions. This is usually expressed as a numerical score that can be calculated through various techniques, such as Shapley Additive Explanations (SHAP) and integrated gradients. For more information, see [Machine learning model interpretability with :AWS](#).

feature transformation

To optimize data for the ML process, including enriching data with additional sources, scaling values, or extracting multiple sets of information from a single data field. This enables the ML model to benefit from the data. For example, if you break down the "2021-05-27 00:15:37" date into "2021", "May", "Thu", and "15", you can help the learning algorithm learn nuanced patterns associated with different data components.

FGAC

See [fine-grained access control](#).

fine-grained access control (FGAC)

The use of multiple conditions to allow or deny an access request.

flash-cut migration

A database migration method that uses continuous data replication through [change data capture](#) to migrate data in the shortest time possible, instead of using a phased approach. The objective is to keep downtime to a minimum.

G

geo blocking

See [geographic restrictions](#).

geographic restrictions (geo blocking)

In Amazon CloudFront, an option to prevent users in specific countries from accessing content distributions. You can use an allow list or block list to specify approved and banned countries. For more information, see [Restricting the geographic distribution of your content](#) in the CloudFront documentation.

Gitflow workflow

An approach in which lower and upper environments use different branches in a source code repository. The Gitflow workflow is considered legacy, and the [trunk-based workflow](#) is the modern, preferred approach.

greenfield strategy

The absence of existing infrastructure in a new environment. When adopting a greenfield strategy for a system architecture, you can select all new technologies without the restriction of compatibility with existing infrastructure, also known as [brownfield](#). If you are expanding the existing infrastructure, you might blend brownfield and greenfield strategies.

guardrail

A high-level rule that helps govern resources, policies, and compliance across organizational units (OUs). *Preventive guardrails* enforce policies to ensure alignment to compliance standards. They are implemented by using service control policies and IAM permissions boundaries. *Detective guardrails* detect policy violations and compliance issues, and generate alerts

for remediation. They are implemented by using AWS Config, AWS Security Hub, Amazon GuardDuty, AWS Trusted Advisor, Amazon Inspector, and custom AWS Lambda checks.

H

HA

See [high availability](#).

heterogeneous database migration

Migrating your source database to a target database that uses a different database engine (for example, Oracle to Amazon Aurora). Heterogeneous migration is typically part of a re-architecting effort, and converting the schema can be a complex task. [AWS provides AWS SCT](#) that helps with schema conversions.

high availability (HA)

The ability of a workload to operate continuously, without intervention, in the event of challenges or disasters. HA systems are designed to automatically fail over, consistently deliver high-quality performance, and handle different loads and failures with minimal performance impact.

historian modernization

An approach used to modernize and upgrade operational technology (OT) systems to better serve the needs of the manufacturing industry. A *historian* is a type of database that is used to collect and store data from various sources in a factory.

homogeneous database migration

Migrating your source database to a target database that shares the same database engine (for example, Microsoft SQL Server to Amazon RDS for SQL Server). Homogeneous migration is typically part of a rehosting or replatforming effort. You can use native database utilities to migrate the schema.

hot data

Data that is frequently accessed, such as real-time data or recent translational data. This data typically requires a high-performance storage tier or class to provide fast query responses.

hotfix

An urgent fix for a critical issue in a production environment. Due to its urgency, a hotfix is usually made outside of the typical DevOps release workflow.

hypercare period

Immediately following cutover, the period of time when a migration team manages and monitors the migrated applications in the cloud in order to address any issues. Typically, this period is 1–4 days in length. At the end of the hypercare period, the migration team typically transfers responsibility for the applications to the cloud operations team.

I

laC

See [infrastructure as code](#).

identity-based policy

A policy attached to one or more IAM principals that defines their permissions within the AWS Cloud environment.

idle application

An application that has an average CPU and memory usage between 5 and 20 percent over a period of 90 days. In a migration project, it is common to retire these applications or retain them on premises.

IIoT

See [industrial Internet of Things](#).

immutable infrastructure

A model that deploys new infrastructure for production workloads instead of updating, patching, or modifying the existing infrastructure. Immutable infrastructures are inherently more consistent, reliable, and predictable than [mutable infrastructure](#). For more information, see the [Deploy using immutable infrastructure](#) best practice in the AWS Well-Architected Framework.

inbound (ingress) VPC

In an AWS multi-account architecture, a VPC that accepts, inspects, and routes network connections from outside an application. The [AWS Security Reference Architecture](#) recommends

setting up your Network account with inbound, outbound, and inspection VPCs to protect the two-way interface between your application and the broader internet.

incremental migration

A cutover strategy in which you migrate your application in small parts instead of performing a single, full cutover. For example, you might move only a few microservices or users to the new system initially. After you verify that everything is working properly, you can incrementally move additional microservices or users until you can decommission your legacy system. This strategy reduces the risks associated with large migrations.

infrastructure

All of the resources and assets contained within an application's environment.

infrastructure as code (IaC)

The process of provisioning and managing an application's infrastructure through a set of configuration files. IaC is designed to help you centralize infrastructure management, standardize resources, and scale quickly so that new environments are repeatable, reliable, and consistent.

industrial Internet of Things (IIoT)

The use of internet-connected sensors and devices in the industrial sectors, such as manufacturing, energy, automotive, healthcare, life sciences, and agriculture. For more information, see [Building an industrial Internet of Things \(IIoT\) digital transformation strategy](#).

inspection VPC

In an AWS multi-account architecture, a centralized VPC that manages inspections of network traffic between VPCs (in the same or different AWS Regions), the internet, and on-premises networks. The [AWS Security Reference Architecture](#) recommends setting up your Network account with inbound, outbound, and inspection VPCs to protect the two-way interface between your application and the broader internet.

Internet of Things (IoT)

The network of connected physical objects with embedded sensors or processors that communicate with other devices and systems through the internet or over a local communication network. For more information, see [What is IoT?](#)

interpretability

A characteristic of a machine learning model that describes the degree to which a human can understand how the model's predictions depend on its inputs. For more information, see [Machine learning model interpretability with AWS](#).

IoT

See [Internet of Things](#).

IT information library (ITIL)

A set of best practices for delivering IT services and aligning these services with business requirements. ITIL provides the foundation for ITSM.

IT service management (ITSM)

Activities associated with designing, implementing, managing, and supporting IT services for an organization. For information about integrating cloud operations with ITSM tools, see the [operations integration guide](#).

ITIL

See [IT information library](#).

ITSM

See [IT service management](#).

L

label-based access control (LBAC)

An implementation of mandatory access control (MAC) where the users and the data itself are each explicitly assigned a security label value. The intersection between the user security label and data security label determines which rows and columns can be seen by the user.

landing zone

A landing zone is a well-architected, multi-account AWS environment that is scalable and secure. This is a starting point from which your organizations can quickly launch and deploy workloads and applications with confidence in their security and infrastructure environment. For more information about landing zones, see [Setting up a secure and scalable multi-account AWS environment](#).

large migration

A migration of 300 or more servers.

LBAC

See [label-based access control](#).

least privilege

The security best practice of granting the minimum permissions required to perform a task. For more information, see [Apply least-privilege permissions](#) in the IAM documentation.

lift and shift

See [7 Rs](#).

little-endian system

A system that stores the least significant byte first. See also [endianness](#).

lower environments

See [environment](#).

M

machine learning (ML)

A type of artificial intelligence that uses algorithms and techniques for pattern recognition and learning. ML analyzes and learns from recorded data, such as Internet of Things (IoT) data, to generate a statistical model based on patterns. For more information, see [Machine Learning](#).

main branch

See [branch](#).

managed services

AWS services for which AWS operates the infrastructure layer, the operating system, and platforms, and you access the endpoints to store and retrieve data. Amazon Simple Storage Service (Amazon S3) and Amazon DynamoDB are examples of managed services. These are also known as *abstracted services*.

MAP

See [Migration Acceleration Program](#).

mechanism

A complete process in which you create a tool, drive adoption of the tool, and then inspect the results in order to make adjustments. A mechanism is a cycle that reinforces and improves itself as it operates. For more information, see [Building mechanisms](#) in the AWS Well-Architected Framework.

member account

All AWS accounts other than the management account that are part of an organization in AWS Organizations. An account can be a member of only one organization at a time.

microservice

A small, independent service that communicates over well-defined APIs and is typically owned by small, self-contained teams. For example, an insurance system might include microservices that map to business capabilities, such as sales or marketing, or subdomains, such as purchasing, claims, or analytics. The benefits of microservices include agility, flexible scaling, easy deployment, reusable code, and resilience. For more information, see [Integrating microservices by using AWS serverless services](#).

microservices architecture

An approach to building an application with independent components that run each application process as a microservice. These microservices communicate through a well-defined interface by using lightweight APIs. Each microservice in this architecture can be updated, deployed, and scaled to meet demand for specific functions of an application. For more information, see [Implementing microservices on AWS](#).

Migration Acceleration Program (MAP)

An AWS program that provides consulting support, training, and services to help organizations build a strong operational foundation for moving to the cloud, and to help offset the initial cost of migrations. MAP includes a migration methodology for executing legacy migrations in a methodical way and a set of tools to automate and accelerate common migration scenarios.

migration at scale

The process of moving the majority of the application portfolio to the cloud in waves, with more applications moved at a faster rate in each wave. This phase uses the best practices and lessons learned from the earlier phases to implement a *migration factory* of teams, tools, and

processes to streamline the migration of workloads through automation and agile delivery. This is the third phase of the [AWS migration strategy](#).

migration factory

Cross-functional teams that streamline the migration of workloads through automated, agile approaches. Migration factory teams typically include operations, business analysts and owners, migration engineers, developers, and DevOps professionals working in sprints. Between 20 and 50 percent of an enterprise application portfolio consists of repeated patterns that can be optimized by a factory approach. For more information, see the [discussion of migration factories](#) and the [Cloud Migration Factory guide](#) in this content set.

migration metadata

The information about the application and server that is needed to complete the migration. Each migration pattern requires a different set of migration metadata. Examples of migration metadata include the target subnet, security group, and AWS account.

migration pattern

A repeatable migration task that details the migration strategy, the migration destination, and the migration application or service used. Example: Rehost migration to Amazon EC2 with AWS Application Migration Service.

Migration Portfolio Assessment (MPA)

An online tool that provides information for validating the business case for migrating to the AWS Cloud. MPA provides detailed portfolio assessment (server right-sizing, pricing, TCO comparisons, migration cost analysis) as well as migration planning (application data analysis and data collection, application grouping, migration prioritization, and wave planning). The [MPA tool](#) (requires login) is available free of charge to all AWS consultants and APN Partner consultants.

Migration Readiness Assessment (MRA)

The process of gaining insights about an organization's cloud readiness status, identifying strengths and weaknesses, and building an action plan to close identified gaps, using the AWS CAF. For more information, see the [migration readiness guide](#). MRA is the first phase of the [AWS migration strategy](#).

migration strategy

The approach used to migrate a workload to the AWS Cloud. For more information, see the [7 Rs](#) entry in this glossary and see [Mobilize your organization to accelerate large-scale migrations](#).

ML

See [machine learning](#).

MPA

See [Migration Portfolio Assessment](#).

modernization

Transforming an outdated (legacy or monolithic) application and its infrastructure into an agile, elastic, and highly available system in the cloud to reduce costs, gain efficiencies, and take advantage of innovations. For more information, see [Strategy for modernizing applications in the AWS Cloud](#).

modernization readiness assessment

An evaluation that helps determine the modernization readiness of an organization's applications; identifies benefits, risks, and dependencies; and determines how well the organization can support the future state of those applications. The outcome of the assessment is a blueprint of the target architecture, a roadmap that details development phases and milestones for the modernization process, and an action plan for addressing identified gaps. For more information, see [Evaluating modernization readiness for applications in the AWS Cloud](#).

monolithic applications (monoliths)

Applications that run as a single service with tightly coupled processes. Monolithic applications have several drawbacks. If one application feature experiences a spike in demand, the entire architecture must be scaled. Adding or improving a monolithic application's features also becomes more complex when the code base grows. To address these issues, you can use a microservices architecture. For more information, see [Decomposing monoliths into microservices](#).

multiclass classification

A process that helps generate predictions for multiple classes (predicting one of more than two outcomes). For example, an ML model might ask "Is this product a book, car, or phone?" or "Which product category is most interesting to this customer?"

mutable infrastructure

A model that updates and modifies the existing infrastructure for production workloads. For improved consistency, reliability, and predictability, the AWS Well-Architected Framework recommends the use of [immutable infrastructure](#) as a best practice.

O

OAC

See [origin access control](#).

OAI

See [origin access identity](#).

OCM

See [organizational change management](#).

offline migration

A migration method in which the source workload is taken down during the migration process. This method involves extended downtime and is typically used for small, non-critical workloads.

OI

See [operations integration](#).

OLA

See [operational-level agreement](#).

online migration

A migration method in which the source workload is copied to the target system without being taken offline. Applications that are connected to the workload can continue to function during the migration. This method involves zero to minimal downtime and is typically used for critical production workloads.

operational-level agreement (OLA)

An agreement that clarifies what functional IT groups promise to deliver to each other, to support a service-level agreement (SLA).

operational readiness review (ORR)

A checklist of questions and associated best practices that help you understand, evaluate, prevent, or reduce the scope of incidents and possible failures. For more information, see [Operational Readiness Reviews \(ORR\)](#) in the AWS Well-Architected Framework.

operations integration (OI)

The process of modernizing operations in the cloud, which involves readiness planning, automation, and integration. For more information, see the [operations integration guide](#).

organization trail

A trail that's created by AWS CloudTrail that logs all events for all AWS accounts in an organization in AWS Organizations. This trail is created in each AWS account that's part of the organization and tracks the activity in each account. For more information, see [Creating a trail for an organization](#) in the CloudTrail documentation.

organizational change management (OCM)

A framework for managing major, disruptive business transformations from a people, culture, and leadership perspective. OCM helps organizations prepare for, and transition to, new systems and strategies by accelerating change adoption, addressing transitional issues, and driving cultural and organizational changes. In the AWS migration strategy, this framework is called *people acceleration*, because of the speed of change required in cloud adoption projects. For more information, see the [OCM guide](#).

origin access control (OAC)

In CloudFront, an enhanced option for restricting access to secure your Amazon Simple Storage Service (Amazon S3) content. OAC supports all S3 buckets in all AWS Regions, server-side encryption with AWS KMS (SSE-KMS), and dynamic PUT and DELETE requests to the S3 bucket.

origin access identity (OAI)

In CloudFront, an option for restricting access to secure your Amazon S3 content. When you use OAI, CloudFront creates a principal that Amazon S3 can authenticate with. Authenticated principals can access content in an S3 bucket only through a specific CloudFront distribution. See also [OAC](#), which provides more granular and enhanced access control.

ORR

See [operational readiness review](#).

outbound (egress) VPC

In an AWS multi-account architecture, a VPC that handles network connections that are initiated from within an application. The [AWS Security Reference Architecture](#) recommends setting up your Network account with inbound, outbound, and inspection VPCs to protect the two-way interface between your application and the broader internet.

P

permissions boundary

An IAM management policy that is attached to IAM principals to set the maximum permissions that the user or role can have. For more information, see [Permissions boundaries](#) in the IAM documentation.

personally identifiable information (PII)

Information that, when viewed directly or paired with other related data, can be used to reasonably infer the identity of an individual. Examples of PII include names, addresses, and contact information.

PII

See [personally identifiable information](#).

playbook

A set of predefined steps that capture the work associated with migrations, such as delivering core operations functions in the cloud. A playbook can take the form of scripts, automated runbooks, or a summary of processes or steps required to operate your modernized environment.

policy

An object that can define permissions (see [identity-based policy](#)), specify access conditions (see [resource-based policy](#)), or define the maximum permissions for all accounts in an organization in AWS Organizations (see [service control policy](#)).

polyglot persistence

Independently choosing a microservice's data storage technology based on data access patterns and other requirements. If your microservices have the same data storage technology, they can encounter implementation challenges or experience poor performance. Microservices are more easily implemented and achieve better performance and scalability if they use the data store best adapted to their requirements. For more information, see [Enabling data persistence in microservices](#).

portfolio assessment

A process of discovering, analyzing, and prioritizing the application portfolio in order to plan the migration. For more information, see [Evaluating migration readiness](#).

predicate

A query condition that returns true or false, commonly located in a WHERE clause.

predicate pushdown

A database query optimization technique that filters the data in the query before transfer. This reduces the amount of data that must be retrieved and processed from the relational database, and it improves query performance.

preventative control

A security control that is designed to prevent an event from occurring. These controls are a first line of defense to help prevent unauthorized access or unwanted changes to your network. For more information, see [Preventative controls](#) in *Implementing security controls on AWS*.

principal

An entity in AWS that can perform actions and access resources. This entity is typically a root user for an AWS account, an IAM role, or a user. For more information, see *Principal* in [Roles terms and concepts](#) in the IAM documentation.

Privacy by Design

An approach in system engineering that takes privacy into account throughout the whole engineering process.

private hosted zones

A container that holds information about how you want Amazon Route 53 to respond to DNS queries for a domain and its subdomains within one or more VPCs. For more information, see [Working with private hosted zones](#) in the Route 53 documentation.

proactive control

A [security control](#) designed to prevent the deployment of noncompliant resources. These controls scan resources before they are provisioned. If the resource is not compliant with the control, then it isn't provisioned. For more information, see the [Controls reference guide](#) in the AWS Control Tower documentation and see [Proactive controls](#) in *Implementing security controls on AWS*.

production environment

See [environment](#).

pseudonymization

The process of replacing personal identifiers in a dataset with placeholder values. Pseudonymization can help protect personal privacy. Pseudonymized data is still considered to be personal data.

Q

query plan

A series of steps, like instructions, that are used to access the data in a SQL relational database system.

query plan regression

When a database service optimizer chooses a less optimal plan than it did before a given change to the database environment. This can be caused by changes to statistics, constraints, environment settings, query parameter bindings, and updates to the database engine.

R

RACI matrix

See [responsible, accountable, consulted, informed \(RACI\)](#).

ransomware

A malicious software that is designed to block access to a computer system or data until a payment is made.

RASCI matrix

See [responsible, accountable, consulted, informed \(RACI\)](#).

RCAC

See [row and column access control](#).

read replica

A copy of a database that's used for read-only purposes. You can route queries to the read replica to reduce the load on your primary database.

re-architect

See [7 Rs](#).

recovery point objective (RPO)

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

recovery time objective (RTO)

The maximum acceptable delay between the interruption of service and restoration of service.

refactor

See [7 Rs](#).

Region

A collection of AWS resources in a geographic area. Each AWS Region is isolated and independent of the others to provide fault tolerance, stability, and resilience. For more information, see [Managing AWS Regions](#) in *AWS General Reference*.

regression

An ML technique that predicts a numeric value. For example, to solve the problem of "What price will this house sell for?" an ML model could use a linear regression model to predict a house's sale price based on known facts about the house (for example, the square footage).

rehost

See [7 Rs](#).

release

In a deployment process, the act of promoting changes to a production environment.

relocate

See [7 Rs](#).

replatform

See [7 Rs](#).

repurchase

See [7 Rs](#).

resource-based policy

A policy attached to a resource, such as an Amazon S3 bucket, an endpoint, or an encryption key. This type of policy specifies which principals are allowed access, supported actions, and any other conditions that must be met.

responsible, accountable, consulted, informed (RACI) matrix

A matrix that defines the roles and responsibilities for all parties involved in migration activities and cloud operations. The matrix name is derived from the responsibility types defined in the matrix: responsible (R), accountable (A), consulted (C), and informed (I). The support (S) type is optional. If you include support, the matrix is called a *RASCI matrix*, and if you exclude it, it's called a *RACI matrix*.

responsive control

A security control that is designed to drive remediation of adverse events or deviations from your security baseline. For more information, see [Responsive controls](#) in *Implementing security controls on AWS*.

retain

See [7 Rs](#).

retire

See [7 Rs](#).

rotation

The process of periodically updating a [secret](#) to make it more difficult for an attacker to access the credentials.

row and column access control (RCAC)

The use of basic, flexible SQL expressions that have defined access rules. RCAC consists of row permissions and column masks.

RPO

See [recovery point objective](#).

RTO

See [recovery time objective](#).

runbook

A set of manual or automated procedures required to perform a specific task. These are typically built to streamline repetitive operations or procedures with high error rates.

S

SAML 2.0

An open standard that many identity providers (IdPs) use. This feature enables federated single sign-on (SSO), so users can log into the AWS Management Console or call the AWS API operations without you having to create user in IAM for everyone in your organization. For more information about SAML 2.0-based federation, see [About SAML 2.0-based federation](#) in the IAM documentation.

SCP

See [service control policy](#).

secret

In AWS Secrets Manager, confidential or restricted information, such as a password or user credentials, that you store in encrypted form. It consists of the secret value and its metadata. The secret value can be binary, a single string, or multiple strings. For more information, see [Secret](#) in the Secrets Manager documentation.

security control

A technical or administrative guardrail that prevents, detects, or reduces the ability of a threat actor to exploit a security vulnerability. There are four primary types of security controls: [preventative](#), [detective](#), [responsive](#), and [proactive](#).

security hardening

The process of reducing the attack surface to make it more resistant to attacks. This can include actions such as removing resources that are no longer needed, implementing the security best practice of granting least privilege, or deactivating unnecessary features in configuration files.

security information and event management (SIEM) system

Tools and services that combine security information management (SIM) and security event management (SEM) systems. A SIEM system collects, monitors, and analyzes data from servers,

networks, devices, and other sources to detect threats and security breaches, and to generate alerts.

security response automation

A predefined and programmed action that is designed to automatically respond to or remediate a security event. These automations serve as [detective](#) or [responsive](#) security controls that help you implement AWS security best practices. Examples of automated response actions include modifying a VPC security group, patching an Amazon EC2 instance, or rotating credentials.

server-side encryption

Encryption of data at its destination, by the AWS service that receives it.

service control policy (SCP)

A policy that provides centralized control over permissions for all accounts in an organization in AWS Organizations. SCPs define guardrails or set limits on actions that an administrator can delegate to users or roles. You can use SCPs as allow lists or deny lists, to specify which services or actions are permitted or prohibited. For more information, see [Service control policies](#) in the AWS Organizations documentation.

service endpoint

The URL of the entry point for an AWS service. You can use the endpoint to connect programmatically to the target service. For more information, see [AWS service endpoints](#) in *AWS General Reference*.

service-level agreement (SLA)

An agreement that clarifies what an IT team promises to deliver to their customers, such as service uptime and performance.

service-level indicator (SLI)

A measurement of a performance aspect of a service, such as its error rate, availability, or throughput.

service-level objective (SLO)

A target metric that represents the health of a service, as measured by a [service-level indicator](#).

shared responsibility model

A model describing the responsibility you share with AWS for cloud security and compliance. AWS is responsible for security *of* the cloud, whereas you are responsible for security *in* the cloud. For more information, see [Shared responsibility model](#).

SIEM

See [security information and event management system](#).

single point of failure (SPOF)

A failure in a single, critical component of an application that can disrupt the system.

SLA

See [service-level agreement](#).

SLI

See [service-level indicator](#).

SLO

See [service-level objective](#).

split-and-seed model

A pattern for scaling and accelerating modernization projects. As new features and product releases are defined, the core team splits up to create new product teams. This helps scale your organization's capabilities and services, improves developer productivity, and supports rapid innovation. For more information, see [Phased approach to modernizing applications in the AWS Cloud](#).

SPOF

See [single point of failure](#).

star schema

A database organizational structure that uses one large fact table to store transactional or measured data and uses one or more smaller dimensional tables to store data attributes. This structure is designed for use in a [data warehouse](#) or for business intelligence purposes.

strangler fig pattern

An approach to modernizing monolithic systems by incrementally rewriting and replacing system functionality until the legacy system can be decommissioned. This pattern uses the analogy of a fig vine that grows into an established tree and eventually overcomes and replaces its host. The pattern was [introduced by Martin Fowler](#) as a way to manage risk when rewriting monolithic systems. For an example of how to apply this pattern, see [Modernizing legacy](#)

[Microsoft ASP.NET \(ASMX\) web services incrementally by using containers and Amazon API Gateway.](#)

subnet

A range of IP addresses in your VPC. A subnet must reside in a single Availability Zone.

symmetric encryption

An encryption algorithm that uses the same key to encrypt and decrypt the data.

synthetic testing

Testing a system in a way that simulates user interactions to detect potential issues or to monitor performance. You can use [Amazon CloudWatch Synthetics](#) to create these tests.

T

tags

Key-value pairs that act as metadata for organizing your AWS resources. Tags can help you manage, identify, organize, search for, and filter resources. For more information, see [Tagging your AWS resources](#).

target variable

The value that you are trying to predict in supervised ML. This is also referred to as an *outcome variable*. For example, in a manufacturing setting the target variable could be a product defect.

task list

A tool that is used to track progress through a runbook. A task list contains an overview of the runbook and a list of general tasks to be completed. For each general task, it includes the estimated amount of time required, the owner, and the progress.

test environment

See [environment](#).

training

To provide data for your ML model to learn from. The training data must contain the correct answer. The learning algorithm finds patterns in the training data that map the input data attributes to the target (the answer that you want to predict). It outputs an ML model that

captures these patterns. You can then use the ML model to make predictions on new data for which you don't know the target.

transit gateway

A network transit hub that you can use to interconnect your VPCs and on-premises networks. For more information, see [What is a transit gateway](#) in the AWS Transit Gateway documentation.

trunk-based workflow

An approach in which developers build and test features locally in a feature branch and then merge those changes into the main branch. The main branch is then built to the development, preproduction, and production environments, sequentially.

trusted access

Granting permissions to a service that you specify to perform tasks in your organization in AWS Organizations and in its accounts on your behalf. The trusted service creates a service-linked role in each account, when that role is needed, to perform management tasks for you. For more information, see [Using AWS Organizations with other AWS services](#) in the AWS Organizations documentation.

tuning

To change aspects of your training process to improve the ML model's accuracy. For example, you can train the ML model by generating a labeling set, adding labels, and then repeating these steps several times under different settings to optimize the model.

two-pizza team

A small DevOps team that you can feed with two pizzas. A two-pizza team size ensures the best possible opportunity for collaboration in software development.

U

uncertainty

A concept that refers to imprecise, incomplete, or unknown information that can undermine the reliability of predictive ML models. There are two types of uncertainty: *Epistemic uncertainty* is caused by limited, incomplete data, whereas *aleatoric uncertainty* is caused by the noise and randomness inherent in the data. For more information, see the [Quantifying uncertainty in deep learning systems](#) guide.

undifferentiated tasks

Also known as *heavy lifting*, work that is necessary to create and operate an application but that doesn't provide direct value to the end user or provide competitive advantage. Examples of undifferentiated tasks include procurement, maintenance, and capacity planning.

upper environments

See [environment](#).

V

vacuuming

A database maintenance operation that involves cleaning up after incremental updates to reclaim storage and improve performance.

version control

Processes and tools that track changes, such as changes to source code in a repository.

VPC peering

A connection between two VPCs that allows you to route traffic by using private IP addresses. For more information, see [What is VPC peering](#) in the Amazon VPC documentation.

vulnerability

A software or hardware flaw that compromises the security of the system.

W

warm cache

A buffer cache that contains current, relevant data that is frequently accessed. The database instance can read from the buffer cache, which is faster than reading from the main memory or disk.

warm data

Data that is infrequently accessed. When querying this kind of data, moderately slow queries are typically acceptable.

window function

A SQL function that performs a calculation on a group of rows that relate in some way to the current record. Window functions are useful for processing tasks, such as calculating a moving average or accessing the value of rows based on the relative position of the current row.

workload

A collection of resources and code that delivers business value, such as a customer-facing application or backend process.

workstream

Functional groups in a migration project that are responsible for a specific set of tasks. Each workstream is independent but supports the other workstreams in the project. For example, the portfolio workstream is responsible for prioritizing applications, wave planning, and collecting migration metadata. The portfolio workstream delivers these assets to the migration workstream, which then migrates the servers and applications.

WORM

See [write once, read many](#).

WQF

See [AWS Workload Qualification Framework](#).

write once, read many (WORM)

A storage model that writes data a single time and prevents the data from being deleted or modified. Authorized users can read the data as many times as needed, but they cannot change it. This data storage infrastructure is considered [immutable](#).

Z

zero-day exploit

An attack, typically malware, that takes advantage of a [zero-day vulnerability](#).

zero-day vulnerability

An unmitigated flaw or vulnerability in a production system. Threat actors can use this type of vulnerability to attack the system. Developers frequently become aware of the vulnerability as a result of the attack.

zombie application

An application that has an average CPU and memory usage below 5 percent. In a migration project, it is common to retire these applications.