
AWS Prescriptive Guidance

Foundation playbook for AWS large migrations



AWS Prescriptive Guidance: Foundation playbook for AWS large migrations

Copyright © 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
About this playbook	1
About the tools and templates	2
People foundation	3
Workstreams	3
Core workstreams	3
Supporting workstreams	8
Roles	12
Team organization	14
Best practices for team organization and composition	14
Creating RACI matrices	16
Cloud Enablement Engine (CEE)	18
Training and skills required	20
Prerequisites	21
Fundamentals	21
Advanced training	22
Create your training dashboard	22
Platform foundation	24
Landing zone considerations	24
Infrastructure considerations	25
Operations considerations	28
Security considerations	30
On-premises considerations	31
Infrastructure considerations	31
Operations considerations	32
Security considerations	32
Document migration principles	33
Resources	36
AWS large migration	36
Strategy	36
Guide	36
Playbooks	36
Additional references	36
AWS Prescriptive Guidance glossary	37
Contributors	45
Document history	46

Foundation playbook for AWS large migrations

Amazon Web Services (AWS)

February 2022 (last update (p. 46): May 2022)

A large migration project is built upon its people foundation and platform foundation. Properly preparing these foundations is critical to the success of the project. *Platform* refers to the technology decisions you make, such as infrastructure, operations, and security. *People* refers to the teams and individuals who contribute to the project, from beginning to end.

In this playbook, you build the foundation workstream. Because this workstream is intended to prepare the platform and people before you begin migrating applications, you start and complete this workstream within the first stage of a large migration, initialization.

About this playbook

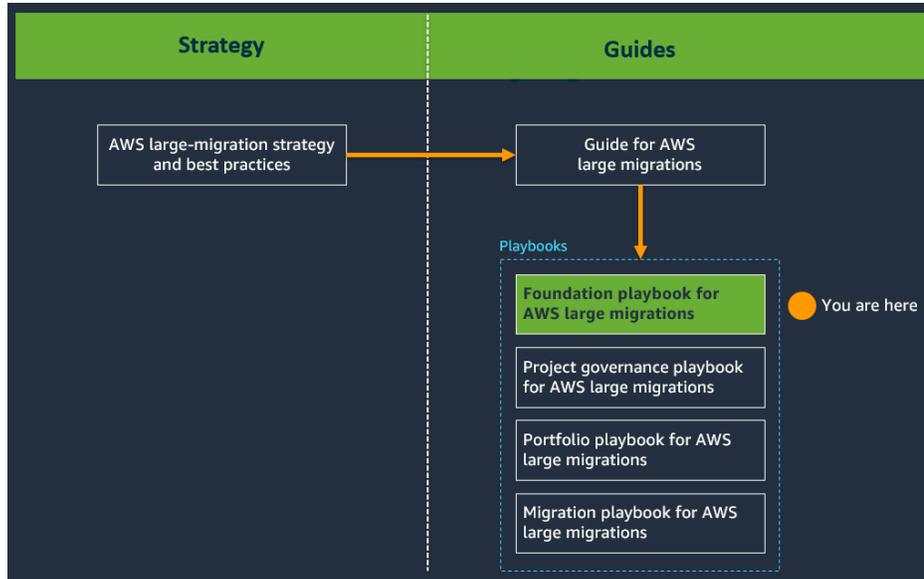
This guide is part of a series about large migrations to the Amazon Web Services (AWS) Cloud. Migrating 300 or more servers is considered a *large migration*. If you haven't already done so, we highly recommend reading the following:

1. [AWS large-migration strategy and best practices](#) – This strategy discusses best practices for large migrations and provides use cases from customers across various industries.
2. [Guide for AWS large migrations](#) – This guide describes a high-level, phased approach to implementing the best practices outlined in the strategy document.

After reading the strategy and guide, the playbooks help you focus on how to implement the strategy and connect all the moving parts of a large migration. A *playbook* helps you build the assets and processes that you use in a large migration, such as runbooks, governing documents, or tools. A *runbook* is a standard operating procedure for a task that you perform repeatedly. The following playbooks are available in this document series for large migrations, and we recommend that you read them in the following order:

1. Foundation playbook for AWS large migrations (this guide)
2. [Project governance playbook for AWS large migrations](#)
3. [Portfolio playbook for AWS large migrations](#)
4. [Migration playbook for AWS large migrations](#)

The following figure shows the structure of the AWS documentation series for large migrations. Review the strategy first, then the guide, and then read the playbooks.



The purpose of this playbook is to prepare the platform foundation and people foundation to support a large-scale migration effort. Both of these foundations are critical to the success of large migrations. This guide consists of the following sections:

1. **Platform foundation** – In this section, you review technology considerations for the on-premises and AWS Cloud environments, such as infrastructure, operations, security. You make key decisions in these categories, which you record as migration principles.
2. **People foundation** – In this section, you define the workstreams in your large migration project and build a responsible, accountable, consulted, informed (RACI) matrix for each high-level task. It also includes recommendations for establishing a Cloud Enable Engine (CEE). This section also contains training resources and helps you build a training dashboard for your large migration.

About the tools and templates

In this playbook, you create the following tools, which you use to prepare the platform and people:

- Migration principles
- RACI matrices
- Dashboard for training

We recommend using the [foundation playbook templates](#) included in this playbook and then customizing them for your portfolio, processes, and environment. The instructions in this playbook tell you when and how to customize each of these templates. This playbook includes the following templates:

- **Dashboard template for training** – This dashboard template helps you build a training plan for each workstream and track each individual's progress toward completing the required training.
- **Data replication calculator** – This workbook helps you estimate the amount of time needed to complete data replication.
- **Migration principles template** – This template helps you record the key infrastructure, operations, and security decisions that you need to make when preparing your platform.
- **RACI template** – This template helps you build a high-level and detailed RACI matrices that outline the roles and responsibilities of your large migration project.

People foundation

This section focuses on preparing the people and processes involved in your project for the activities in each stage of the large migration. To build the people foundation, you need to define the workstreams in your project, organize individuals into functional teams, confirm that that roles and responsibilities are well understood, and complete training.

This section consists of the following topics:

- [Workstreams in a large migration \(p. 3\)](#)
- [Roles \(p. 12\)](#)
- [Team organization and composition \(p. 14\)](#)
- [Training and skills required for large migrations \(p. 20\)](#)

Workstreams in a large migration

Large migration projects typically consist of multiple workstreams, and each workstream has a clear scope of tasks. Each workstream is independent but also supports the other workstreams to accomplish the same goal – migrate servers at scale. This section discusses the standard core workstreams for large migrations as well as common supporting workstreams.

Core workstreams

Core workstreams are needed for every large migration, regardless of company size or segment. The following is an overview of the primary roles of each core workstream:

- **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.

Information and activities flow from upstream to downstream in a large migration, as shown in the following table. Information comes from the upstream foundation and project governance workstreams, through the portfolio workstream, and into the migration workstream. For example, the portfolio workstream is upstream of the migration workstream because the portfolio workstream prepares the metadata and wave plan that the migration workstream uses to migrate and cutover the applications and servers. Adding additional, supporting workstreams in your large migration project might change the flow of information and activities through the core workstreams.

Important

You need to assign a project-level technical leader for your large migration project. This role is not part of any individual workstream but has the total responsibility of all workstreams. This

AWS Prescriptive Guidance Foundation
playbook for AWS large migrations
Core workstreams

individual oversees all workstreams to make sure they work together and stay focused on the project-level goals.

Core workstream name	Upstream workstreams	Downstream workstreams
Foundation	—	Migration Portfolio
Project governance	—	Migration Portfolio
Portfolio	Foundation Project governance	Migration
Migration	Foundation Project governance Portfolio	—

The following are the primary functions of each core workstream in the phases of a large migration. The playbooks in this document series are structured to help you navigate the tasks for each workstream in the appropriate phase and stage.

		Foundation	Project governance	Portfolio	Migration
Phase 1: Assess		—	—	—	—
Phase 2: Mobilize		You might have designed the AWS landing zone or workstreams in this phase.	You might have designed a project management process in this phase.	You might have completed an initial portfolio assessment and discovery in this phase.	You might have completed a pilot migration in this phase.
Phase 3: Migrate	Stage 1: Initialize	Establish workstreams and review landing zone design. Prepare for change. Formalize migration principles, teams, and RACI matrix. Complete training.	Develop project management processes and communication and meeting plans.	Develop metadata, wave planning, and application prioritization runbooks.	Develop migration runbooks.
	Stage 2: Implement	—	Facilitate and communicate	Collect metadata for	Migrate and cutover waves,

		Foundation	Project governance	Portfolio	Migration
			the status of waves and the overall migration project.	the migration, prioritize applications, and plan waves.	and iterate the runbooks to increase velocity.

The following sections describe each of the core workstreams in more detail, including common tasks for each workstream, the expected outcome of each workstream, and the skills required in each workstream. It is not required that each individual in the workstream have every skill. A workstream consists of one more cross-functional teams, so each person contributes different skills. But as a team, they should have all the skills listed.

Foundation workstream

The foundation workstream consists of two categories: platform foundation and people foundation. Building the platform foundation helps confirm that both the AWS and the on-premises infrastructures are ready to support the large migration. Building a people foundation prepares and trains the project teams for the migration and sets up all workstreams.

Common tasks	<ul style="list-style-type: none"> • Build and validate the AWS landing zone • Prepare the on-premises infrastructure to support the migration, such as making networking or firewall changes, permissions changes, or Active Directory changes • Set up the project core workstreams and supporting workstreams • Set up the training plan for the team • Build the RACI matrices with project managers
Expected outcome	<ul style="list-style-type: none"> • Source and target platforms are prepared for the large migration. • People are ready to support the large migration • All workstreams are set up.
Required skills	<ul style="list-style-type: none"> • Deep knowledge of on-premises data centers, including servers, storage, and networking • Experience with the AWS Cloud and knowledge of AWS compute services, including landing zones and AWS Control Tower • Experience with large data center or cloud migrations • Experience building a training plan • Experience building a cross-functional team

Project governance workstream

The project governance workstream manages the overall migration project and is responsible for delivering the project on budget and on time.

Common tasks	<ul style="list-style-type: none"> • Kick off the project • Set up the governance model • Set up the Cloud Enablement Engine (CEE) • Set up the communication plan • Set up the escalation plan • Build RACI matrices • Set up the project management framework • Set up status reporting and project tracking • Set up risk and issue tracking • Continuously manage the project by using the predefined processes and tools
Expected outcome	<ul style="list-style-type: none"> • Ensure that every workstream is able to complete their tasks on time • Ensure collaboration across workstreams • Ensure that the project achieves the defined business outcomes • Deliver the project on budget and on time
Required skills	<ul style="list-style-type: none"> • Experience with common project management methodologies, such as waterfall, agile, Kanban, and scrum • Experience with common project management tools, such as Jira, Microsoft Project, and Confluence • Experience with large migration project management

Portfolio workstream

The portfolio workstream manages all of the migration discovery activities, collects metadata, prioritizes applications, and creates a wave plan to support the migration workstream.

Common tasks	<ul style="list-style-type: none"> • Validate the migration strategies and patterns • Complete portfolio discovery by using discovery tools and configuration management database (CMDB) • Define the required metadata, collection processes, and storage location • Prioritize applications • Perform application deep dives, including dependency analysis and target state design • Perform wave planning • Collect migration metadata
Expected outcome	<ul style="list-style-type: none"> • Continuously create wave plans and collect migration metadata, and then hand off to the migration workstream

Required skills	<ul style="list-style-type: none">• Deep knowledge of on-premises CMDB, data repositories, and content management tools• Experience with common portfolio discovery tools, such as AWS Application Discovery Service, Flexera One, and modelizeIT• Experience with portfolio assessment and application prioritization• Experience with application deep dives and application owner interviews• Experience with application designs for the AWS Cloud• Experience with wave planning for large migrations• Experience with automation, including shell scripting, Python, and Microsoft PowerShell
-----------------	---

Migration workstream

The migration workstream manages the migration implementation-related activities, including data replication and cutover. Because the migration team performs the migration and cutover, a common misconception is that the migration workstream does everything in a large migration project. However, the migration workstream is dependent on other workstreams to build the foundation and provide portfolio data to support the migration.

Tip

The migration workstream is generally the largest workstream in a large migration project. Depending on the size and strategy of your project, consider dividing this workstream into multiple sub-workstreams. For example:

- Rehost migration workstream
- Replatform migration workstream
- Refactor migration workstream
- Relocate migration workstream
- Migration workstream for a specialized workload, such as SAP or databases

Common tasks	<ul style="list-style-type: none">• Validate the migration wave plans• Build the migration runbooks• Use AWS migration services to transfer data, such as AWS Application Migration Service (AWS MGN), AWS Database Migration Service (AWS DMS), and AWS DataSync• Install and uninstall software on source and target servers as needed support the migration• Write automation scripts to automate migration activities• Launch target AWS environments, such as Amazon Elastic Compute Cloud (Amazon EC2) instances, for testing or cutover• Work with change management team for changes and cutovers
--------------	---

	<ul style="list-style-type: none"> • Perform migration cutover • Support application owners during application testing • If cutover fails, help roll back the server
Expected outcome	<ul style="list-style-type: none"> • Complete migration cutover and application go-live in target AWS accounts
Required skills	<ul style="list-style-type: none"> • Deep knowledge of on-premises data centers, including servers, storage, and networking • Experience with the AWS Cloud and knowledge of AWS compute services, including landing zone and AWS Control Tower • Experience with AWS migration services, including Application Migration Service, AWS DMS, DataSync, and AWS Snow Family • Experience with large data center or cloud migrations and cutovers • Experience with automation, including shell scripting, Python, and Microsoft PowerShell

Supporting workstreams

Supporting workstreams support the core workstreams. These workstreams are optional, and you might decide to use them based on your use case and the current stage of your migration. The following are some common supporting workstreams that you might want to include in your large migration project:

- **Security and compliance workstream** – This workstream defines and builds the security standards for the target AWS infrastructure and supports migrations.
- **Cloud operations (Cloud Ops) workstream** – This workstream manages applications after cutover, when the hypercare period is complete.
- **Application testing workstream** – This workstream performs application testing before and during the cutover.
- **Specialized workload migration workstream** – This workstream supports migrations for specific, specialized workloads, such as SAP or databases.

You might not need a dedicated workstream for these activities. It is common to have an individual or set of individuals be responsible for these activities and then embed those individuals in one of the core workstreams. For example, every large migration requires a security and compliance person because you need to make sure your target infrastructure is secure and compliant. However, security and compliance assessments and decisions are typically performed early in the migration, most commonly in the mobilize phase. If you have already completed this, you do not need a dedicated workstream to repeat the same tasks. However, it is recommended that you embed a security and compliance person in the migration workstream in order to support the migration activities.

When you add supporting workstreams, it modifies the flow of information and activities through the core workstreams. The following table is an example of how adding workstreams changes this flow. Your supporting workstreams might differ from the examples in this table.

Workstream name	Type	Upstream workstreams	Downstream workstreams
Migration	Core	Foundation	Application testing

Workstream name	Type	Upstream workstreams	Downstream workstreams
		Project governance Portfolio Security and compliance	Cloud operations
Portfolio	Core	Foundation Project governance Security and compliance	Migration
Project governance	Core	—	Migration Portfolio
Foundation	Core	—	Migration Portfolio Cloud operations
Security and compliance	Supporting	—	Migration Portfolio
Cloud operation	Supporting	Migration Application testing Foundation	—
Application testing	Supporting	Migration	Cloud operations
Specialized workload migration	Supporting	Foundation Project governance Portfolio Security and compliance	Application testing Cloud operations

Security and compliance workstream

The security and compliance workstream defines and builds the security standards for AWS infrastructure and supports migrations. Using the standards established by this workstream, application owners typically define the security and compliance requirements for each application. You might decide to have the security and compliance workstream review and approve the requirements for some or all applications.

Common tasks	<ul style="list-style-type: none"> Define the security requirements for the AWS landing zone, such as centralized logging, encryption, AWS Identity and Access
--------------	---

	<ul style="list-style-type: none"> Management (IAM) policies, and Active Directory integration Define the compliance requirements, such as HIPAA, personally identifiable information (PII), Service Organization Control (SOC), and Federal Risk and Authorization Management Program (FedRAMP) Define the security requirements for the migration, such as firewall, security group, and IAM role requirements Manage changes for security-related tasks, such as changes to firewalls, security groups, and permissions
Expected outcome	<ul style="list-style-type: none"> Complete migration cutover and application go-live in target AWS accounts
Required skills	<ul style="list-style-type: none"> Deep knowledge of on-premises data centers, including servers, storage, and networking Deep knowledge of the specialized workload in scope Experience with the AWS Cloud and knowledge of AWS compute services, including landing zones and AWS Control Tower Experience with AWS migration tools, including Application Migration Service, AWS DMS, DataSync, and AWS Snow Family Experience with large data center or cloud migrations and cutovers

Cloud operations workstream

The cloud operations workstream supports the applications after migration cutover. Sometimes cloud operations is in a separate workstream with dedicated resources, but most commonly, these resources come from existing IT operations teams. In that case, no dedicated workstream is required.

Common tasks	<ul style="list-style-type: none"> Monitor and back up the migrated servers and applications Manage the business-as-usual service requests from the application teams, such as increasing the disk size or changing instance types Resolve any application issues and outages as needed Manage the patching policies and schedules Manage the maintenance tasks and requests
Expected outcome	<ul style="list-style-type: none"> Migrated servers and applications are running smoothly on AWS Respond to service requests from users and resolve any issues
Required skills	<ul style="list-style-type: none"> Deep understanding of how the on-premises data center currently operates

	<ul style="list-style-type: none"> • Experience with common AWS operations services, such as Amazon CloudWatch, AWS Config, AWS CloudTrail, AWS Backup, AWS Support • Experience with troubleshooting, and understands the SLA • Experience with supporting large migrations
--	---

Application testing workstream

The application testing workstream supports application testing before and during the cutover. This workstream is more common in projects where system integrators manage the data centers because the application owners don't have sufficient knowledge to perform the application tests. In most cases, the application owner performs these activities, and a dedicated application testing workstream is not required.

Common tasks	<ul style="list-style-type: none"> • Perform application testing before the cutover • Perform application testing during the cutover • Make application changes as needed to work in the new environment • Make a go or no-go decision for applications based on testing results during cutover
Expected outcome	<ul style="list-style-type: none"> • Complete application testing on time during cutover • Make application changes as needed to support the target environment
Required skills	<ul style="list-style-type: none"> • Deep knowledge of the applications and how they operate on premises • Experience with the AWS Cloud, especially the target AWS services • Experience with large migrations

Migration workstream for a specialized workload

You can create a migration workstream that is dedicated to specialized workloads. Generally, you can build standard migration patterns and runbooks to migrate servers and applications at scale, and these are managed by the migration workstream. However, in some cases, certain applications require special migration processes. For example, you might need a special process in order to migrate Hadoop workloads, SAP HANA databases, or mission-critical applications that cannot tolerate the standard amount of down time. For more information about specialized workloads, see *MAP specialized workloads* at [AWS Migration Acceleration Program](#).

Common tasks	<ul style="list-style-type: none"> • Validate the migration wave plans • Build migration runbooks • Use migration tools or native application tools to transfer data • Launch target AWS environments, such as EC2 instances, for testing or cutover
--------------	--

	<ul style="list-style-type: none"> • Work with the change management team for changes and cutovers • Perform migration cutover • Support application owners during application testing • If cutover fails, roll back the application or server
Expected outcome	<ul style="list-style-type: none"> • Complete migration cutover and application go-live in target AWS accounts
Required skills	<ul style="list-style-type: none"> • Deep knowledge of on-premises data centers, including servers, storage, and networking • Deep knowledge of the specialized workload in scope • Experience with the AWS Cloud and knowledge of AWS compute services, including landing zones and AWS Control Tower • Experience with AWS migration tools, including Application Migration Service, AWS DMS, DataSync, and AWS Snow Family • Experience with large data center or cloud migrations and cutovers • Experience with migrating the specialized workload

Roles

The following are the common roles in a large migration project. Because these roles might go by another title in your organization, a brief description of each role is provided. If a role is not available in your organization, you might investigate whether other resources in your organization can perform this role or seek outside support in the form of a consultant.

General role	Alternate titles	Workstreams	Characteristics
Application owner	Application architect, application project coordinator, application project manager	All	Should have in-depth knowledge of their applications
Automation engineer	DevOps engineer	Migration, portfolio	Should have experience and in-depth knowledge of how to build automation scripts
Cloud architect	Cloud engineer, migration consultant, architecture lead, cloud infrastructure architect	Migration, foundation, portfolio	Should have experience and in-depth knowledge of how to design the AWS Cloud infrastructure, how to perform portfolio assessment and wave planning, and how to

AWS Prescriptive Guidance Foundation
playbook for AWS large migrations
Roles

General role	Alternate titles	Workstreams	Characteristics
			use migration tools to migrate workloads to the AWS Cloud
Cloud operations lead	Migration technical support, cloud operations workstream lead	Cloud operations	Should have experience and in-depth knowledge of how to operate workloads in the AWS Cloud
Communication lead	Business unit liaison	Project governance	Should have relationship to the business unit and manage all communications
Executive leadership	Project sponsor	All	Should have clear vision of the migration project
Migration lead	Migration support lead, migration technical product owner, migration workstream lead	Migration	Should have experience and in-depth knowledge of all migration patterns and how to use migration tools to migrate workloads to the AWS Cloud
Portfolio lead	Discovery lead, wave planning lead, portfolio workstream lead	Portfolio	Should have experience and in-depth knowledge of how to perform discovery, portfolio assessment, and wave planning
Project manager	Program manager, project coordinator, Scrum master, project delivery lead, program delivery lead, large migration manager	Project governance	Should have experience and in-depth knowledge of how to manage a large migration project and how to use agile methodologies
Project technical lead	Engineering lead, technical lead, chief architect	All	Should have experience and in-depth knowledge of all workstreams and how to deliver a migration project from start to finish. Responsible for the entire project outcome across all workstreams

General role	Alternate titles	Workstreams	Characteristics
System integrator	Global system integrator	All	Varies, depending on the workstream. Should have in-depth knowledge of workstream-level activities, such as portfolio assessment or server migration
Testing lead	Testing specialist, application testing workstream lead	Application testing	Should have experience and in-depth knowledge of how to perform application testing in the AWS Cloud

Team organization and composition

This section includes the following topics:

- [Best practices for team organization and composition \(p. 14\)](#)
- [Creating RACI matrices \(p. 16\)](#)
- [Cloud Enablement Engine \(CEE\) \(p. 18\)](#)

Best practices for team organization and composition

Team composition in a large migration varies by organization and changes over the course of the project. The following are best practices that are common for all large migration projects:

- **Identify a single, threaded technical leader at the project level and avoid silos** – Large migration projects often have multiple workstreams and teams, each team has different tasks and expected outcomes. A single, threaded leader at the project level is important because this leader makes sure all workstreams work together and stay connected. This helps prevent silos and boundaries. For example, the portfolio workstream needs to continuously send the migration metadata to the migration workstream to support the migration activities. Without a complete understanding of the required migration metadata, the output of the portfolio workstream might not work as an input for the migration workstream. A single, threaded leader helps coordinate the inputs and outputs of each workstream to help the migration run efficiently.
- **Align all workstream-level outcomes with the project-level business outcomes** – Project-level business outcomes should be communicated to all workstream leaders before the migration starts. Each workstream leader must understand the role of their workstream and design their processes to support the project-level business outcomes. For example, if a project-level business outcome is exiting a data center in the next 12 months and speed is the most important factor, the workstream leaders should do the following:
 - All workstreams should prioritize rehost migrations, reduce the number of manual tasks, and add automation to improve the velocity.
 - The portfolio workstream should define standardized patterns and limit customizable patterns to reduce the amount of time required to design the target environment.
- **Design workstreams based on project scope and stage** – Every migration project is different, and one size does not fit all. We recommend having four core workstreams for all large migration projects: migration workstream, portfolio workstream, project governance workstream, and foundation

workstream. You might decide to create additional, supporting workstreams depending on your use case. For more information about workstreams, see [Workstreams in a large migration \(p. 3\)](#). For example, if you have not yet designed the security guardrails in the mobilize phase, you need to create a security and compliance workstream that can define the security and compliance requirements before you start migrating. For more information about building the security guardrails in the mobilize phase, see [Security, risk and compliance](#) in *Mobilize your organization to accelerate large-scale migrations*.

- **Get the application team involved before the migration** – A large migration is never just an IT infrastructure project – it changes the operating model for your business. Involving the application team early and embedding the application owners into your large migration workstreams is critical to the success of large migration project. For example, during portfolio assessment, schedule your meetings early with application owners so that they can participate in the deep dive and help design their application’s target state on AWS.
- **Determine the team size based on workstreams and business outcomes** – Your expected business outcomes and migration strategies drive the size of each team, which is composed of smaller units called pods. In each workstream, you define teams for each migration strategy and then separate those teams into pods. For example, if rehost is your primary migration strategy, then you should have a rehost migration team that is composed of pods that contain 3–5 people. When operating at peak velocity, a pod of 4–5 people on a migration team can typically rehost up to 50 servers per week. This is approximately 200 servers per month or 2,500 servers per year. If your target is to rehost 100 servers per week, you should create two pods of 4–5 people within the rehost migration team. If you are targeting less than 50 per week, you can reduce the size of the migration pod to 3 people. Replatform migrations usually cost more than rehost, and the same size pod can migrate up to 20 servers per week. The portfolio workstream is usually half the size of the migration workstream. You create additional teams and pods in each workstream to support each migration strategy. These recommendations assume that your migration resources are skilled and do not require significant training. The following table is an example of how you would divide the migration and portfolio workstreams into teams and pods for the rehost and replatform migration strategies. The following example assumes that you need to migrate 120 servers per week (100 rehost + 20 replatform) or 6,000 servers per year. This example is the maximum velocity. We recommend that you plan for additional resources in order to help prevent delays.

Workstream	Team	Pod	Resources
Migration workstream	Rehost migration team	Rehost migration pod 1	4–5 people
		Rehost migration pod 2	4–5 people
	Replatform migration team	Replatform migration pod	4–5 people
Portfolio workstream	Portfolio team	Portfolio pod 1	3–4 people
		Portfolio pod 1	3–4 people

- **Build a governance model in the early stage** – A large migration typically involves many people, including people from your own company, third-party software vendors, system integrators, or external consultants. Your project might include representatives from AWS, such as your account team, support engineers, or experts from AWS Professional Services. Your delivery model varies depending on your project scope and who you work with to deliver the project. For example, your project might include AWS or a system integrator, or you might include both. It is important to build a governance model early and create a RACI matrix that clearly defines the roles and responsibilities. As a recommendation, we also recommend creating a Cloud Enablement Engine (CEE), also known as *Cloud Center of Excellence*, in your organization and including representation from all parties. The key purpose of the CEE is to transform the organization from an on-premises operating model to a cloud-operating model. This centralized team is critical to the success of a large migration because it

manages relationships, makes key decisions, and handles escalations throughout the project. The CEE is discussed in more detail later in this guide.

Creating RACI matrices

A large migration project typically involves a lot of people, so building a governance model is important to manage the project. One of the key components of a governance model is a *RACI matrix*, which is used to define the roles and responsibilities for all parties involved in the large migration. The name RACI matrix is derived from the four responsibility types defined in the matrix:

- **Responsible (R)** – This role is responsible for performing the work to complete the task.
- **Accountable (A)** – This role is held accountable for making sure the task is completed. This role is also responsible for ensuring the prerequisites are met and delegating the task to those who are responsible.
- **Consulted (C)** – This role should be consulted for opinions or expertise on the task. Depending on the task, this responsibility type might not be required.
- **Informed (I)** – This role should be kept up to date on the progress of the task and notified when the task is completed.

Because of the complexity of a large migration, we do not recommend using a single RACI matrix to document every task in the large migration. A multi-layer RACI matrix is a much more accessible approach. You start by building a high-level RACI matrix, and then you add more details to each section to build a detailed matrix. Building a detailed RACI matrix is not a one-off approach. You need to build new matrices or add more details to the existing ones as you progress through the portfolio and discover more migration strategies and patterns.

In the [foundation playbook templates](#), you can use the RACI template (Microsoft Excel format) as a starting point for building your own high-level and detailed RACI matrices. This template includes two examples of detailed RACI matrices, one for a rehost migration and another for a replatform migration. The tasks in these examples are included for sample purposes only, and you should customize these examples based on your use case.

Build a high-level RACI matrix

Before you start building a high-level RACI matrix, you need to have the following information ready:

- **Who are the high-level parties involved in this migration?** Identify any partners or consultants that will be involved in this project, such as AWS Professional Services or system integrators. Consider whether any part of your current IT infrastructure is managed by an external partner. The following are examples of high-level parties:
 - Your organization
 - AWS Professional Services
 - System integrators
- **What are the workstreams in your migration?** For more information, see [Workstreams in a large migration \(p. 3\)](#). At a minimum, you should have the four core workstreams, and you can add support workstreams as needed for your project.
- **What are the high-level tasks in your migration?** Create a list of the high-level tasks in your migration. The following are examples of high-level tasks:
 - Build an AWS landing zone
 - Perform portfolio assessment and collect migration metadata
 - Perform a rehost, replatform, or relocate migration
 - Perform application testing and cutover

- Perform project management and governance tasks

Do the following to build your high-level RACI matrix:

1. In the [foundation playbook templates](#), open the *RACI template* (Microsoft Excel format).
2. On the **High-level RACI** tab, in the first row, enter your organization name and any partners that you identified.
3. In the first column, enter the high-level tasks and workstreams that you identified.
4. In the matrix, determine which parties are responsible for each task as follows:
 - If a party is **responsible** for completing the task, enter an **R**.
 - If a party is **accountable** for the task, enter an **A**.
 - If a party should be **consulted** about the task, enter a **C**.
 - If a party should be **informed** about the task, enter an **I**.

The following table is an example of a high-level RACI matrix.

Task	Your organization	Partner A	Partner B	Partner C
Build an AWS landing zone	R/C	A	I	I
Perform portfolio assessment and wave planning	R/C	A	I	I
Perform rehost migration activities	C	C	R/A	I
Perform replatform migration activities	C	C	I	R/A
Project management and governance	R/C	A	I	I
Application changes and testing	C	R/A	C	C
Cloud operations	I	C	R/A	I

Build the detailed RACI matrices

After creating the high-level RACI matrix, the next step is to create a detailed RACI for each high-level task and further refine the tasks, parties, and ownership. Before you start building detailed matrices, you need to have the following information ready:

- **What are the detailed tasks in your migration?** After you have prepared the runbooks and task lists for your large migration project, the processes and details in these runbooks form the detailed layer of your RACI matrix. For example, for a rehost migration, detailed tasks might include installing a

replication agent, verifying replication, and launching test instances for boot-up testing. If you haven't done so already, follow the instructions in the following playbooks to create these documents:

- [Portfolio playbook for AWS large migrations](#)
- [Migration playbook for AWS large migrations](#)
- **What smaller teams make up each workstream and each high-level party?** For example, teams in your organization might include an application team, infrastructure team, operations team, networking team, or a project management office.

Do the following to build a detailed RACI matrix:

1. Open your high-level RACI matrix.
2. Create a copy of the **Detailed RACI (template)** spreadsheet.
3. Name the copied spreadsheet for a high-level task that you identified in [Build a high-level RACI matrix \(p. 16\)](#).
4. In the first row, enter the names of the teams involved in this high-level task.
5. In the first column, enter the detailed tasks that you identified for this high-level task. You can group the detailed tasks into logical sequential groups, which helps readers navigate the matrix.
6. In the matrix, determine which teams are responsible for each task as follows:
 - If a team is **responsible** for completing the task, enter an *R*.
 - If a team is **accountable** for completing the task, enter an *A*.
 - If a team should be **consulted** about the task, enter a *C*.
 - If a team should be **informed** about the task, enter an *I*.
7. For each detailed task, confirm that only one team is responsible and only one team is accountable. If multiple teams are responsible or accountable, this can indicate that the task is not clearly defined or doesn't have clear ownership.
8. Share the detailed RACI matrix with the identified teams and confirm that all teams are familiar with their roles and responsibilities.
9. Repeat this process for each high-level task that you identified in [Build a high-level RACI matrix \(p. 16\)](#).

For examples of detailed RACI matrices, see the **Rehost RACI** and **Replatform RACI** spreadsheets in the *RACI template*, available in the [foundation playbook attachments](#).

Cloud Enablement Engine (CEE)

Best practices for using a CEE

The purpose of a CEE is transforming an IT organization from an on-premises operating model to a cloud-operating model, and it is responsible for guiding the organization through the organizational and cultural changes. As a best practice, it is recommended that you establish a CEE for your large migration. The well-defined foundational processes and guard rails of a CEE can help you achieve the scale and velocity required for large migrations. For information about setting up a CEE, see [Cloud Enablement Engine: A Practical Guide](#). The following are additional recommendations and best practices for establishing a CEE for a large migration project:

- The CEE team should be comprised of cross-functional leaders with the following qualities:
 - Have deep institutional knowledge
 - Have strong, long-standing internal relationships
 - Have a vested interest in the progress and success in the large migration
 - Are curious and want to learn

- Are primarily or solely focused on the migration
- The CEE team should be a mix of people who have worked together previously and newcomers who can provide fresh insights.
- The CEE team should have strong executive support and alignment on the migration objectives.
- Make sure the goals of the CEE team are specific to the large migration.
- Conduct regular, open meetings that provide opportunities for questions and answers, demonstrate cloud services and architectures, and share updates on successful migrations and other wins.
- The CEE team should be empowered to make critical decisions about the large migration project.

Typical CEE roles and responsibilities for large migrations

The following table provides roles in a large-migration CEE team, and it describes the typical tasks and responsibilities for each role. The actual composition of your team and their responsibilities can vary based your use case, scope, and business objective.

Roles	Tasks and responsibilities
Executive sponsor	<ul style="list-style-type: none"> • Managing escalations • Aligning the organization tightly around the objectives and criticality of the migration. • Serving as the voice of authority
Enterprise architect or project-level technical lead	<ul style="list-style-type: none"> • Identifying and documenting the reference architecture for known workload types • Designing and building migration processes for the entire project, across all workstreams • Serving as the single, threaded technical leader who makes sure all workstreams are collaborating and working to deliver the same business-level objectives • Strong institutional knowledge of major applications and common architectures
Project management office lead	<ul style="list-style-type: none"> • Managing timelines, onboarding, training, documentation, reporting, communication, and resource governance • Managing resourcing and training • Managing migration-related town halls
Migration lead	<ul style="list-style-type: none"> • Designing migration processes and tools • Designing migration strategies and automation • Overseeing migration cutovers and achieving the target velocity
Portfolio lead	<ul style="list-style-type: none"> • Designing portfolio assessment and wave planning processes and tools • Designing portfolio discovery and data collection processes • Overseeing the continuous supply of migration metadata and wave plans
Cloud operations lead	<ul style="list-style-type: none"> • Designing the operating model for running workloads on AWS

Roles	Tasks and responsibilities
	<ul style="list-style-type: none"> • Designing strategies for monitoring, incident response, tagging, business continuity, and disaster recovery strategies
Application team leader	<ul style="list-style-type: none"> • Managing the relationship with individual application owners • Managing migration planning and cutovers for their applications • Managing application changes, testing, and approvals
Network and infrastructure lead	<ul style="list-style-type: none"> • Designing the AWS landing zone for target accounts • Designing network connectivity and infrastructure • Designing and deploying security groups • Managing infrastructure and networking changes to support the large migration
Licensing lead	<ul style="list-style-type: none"> • Identifying all commercial off-the-shelf (COTS) and enterprise applications and working with the migration team and application team to plan migration strategies around licensing
Security and compliance lead	<ul style="list-style-type: none"> • Designing authentication and authorization for the large migration, including Active Directory, single sign-on, and IAM policies • Designing network security, including on-premises firewalls, and managing vulnerabilities • Designing compliance requirements for in-scope workloads

Training and skills required for large migrations

The people involved in the large migration are a critical resource, and it is equally as important to prepare them for the migration as it is to prepare the landing zone or workstreams. This section is dedicated to training the people in your project, ensuring that your teams have the skills necessary to perform a large migration. While some skills are common and required for many roles, other skills are more specialized and require thoughtful recruitment or training. By ensuring individuals are properly trained for their roles before the migration starts, the workstreams can operate efficiently, and you can quickly ramp up the migration to the target velocity.

Training is divided into levels: prerequisites, fundamentals, and advanced. Every person in your large migration project should complete the prerequisite-level training, which reviews basic information about the AWS Cloud and migration concepts. For fundamentals and advanced levels, you use a training plan to assign a training level to each workstream. You then use a training tracking tool to record each individual's progress toward completing the required trainings in their workstream. It is important to note that we recommend training based on workstreams rather than roles and job titles because roles can vary significantly between organizations.

Each of the following sections lists and describes the training resources recommended for the level:

- [Large migration training – Prerequisites \(p. 21\)](#)

- [Large migration training – Fundamentals \(p. 21\)](#)
- [Large migration training – Advanced \(p. 22\)](#)

Prerequisites

At a minimum, the resources in every workstream should have foundational understanding of infrastructure, networking, and core AWS services, AWS Cloud Adoption Framework (AWS CAF) and the Well-Architected Framework. The following are recommended for this training level:

- **[AWS Technical Essentials](#)** – This foundational training module provides an overview of AWS services and cloud technology, such as virtual private clouds (VPCs), Amazon Elastic Compute Cloud (Amazon EC2), Availability Zones, and AWS Regions.
- **Foundational training for infrastructure, networking, and data centers** – Provide foundational training about infrastructure and networking, such as Transmission Control Protocol (TCP), Internet Protocol (IP), Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), and load balancers. Provide training about data center technologies, such as the software development lifecycle (SDLC) and IT service management (ITSM). Training requirements in this category vary based on your environment and use case, and many training resources are available. We recommend working with your IT department to identify technology-level training that is appropriate for all personnel in your large migration project
- **Organizational processes** – Provide training for any processes that are specific to your organization, such as change management processes. You must understand the deadlines, approvals, and formal documents required to make changes in your organization, such as firewall and domain changes. Determine whether external partners or consultants need this training in order to support your project.
- **Shared Responsibility Model** – If you are working with AWS Professional Services, this webpage describes how you will share roles and responsibilities with AWS.
- **An Overview of the AWS Cloud Adoption Framework (AWS CAF)** – This whitepaper helps you understand the goals of AWS CAF, the AWS CAF perspective, and the stakeholders involved.

Fundamentals

This section provides an overview of the processes, tools, and guidelines required to successfully complete a large migration. The following are recommended for this training level:

- **How to migrate** This webpage helps you understand the three-phase migration process.
- **About the migration strategies** – This section of the *Guide for AWS large migrations* describes each of the migration strategies and common use cases for each in a large migration project.
- **Migrating to AWS: A high level introduction** – This course provides an overview of the key topics and target audience of the *Migrating to AWS* classroom course.
- **Migrating to AWS** – This course explains how to plan and migrate existing workloads to the AWS Cloud.
- **AWS large-migration strategy and best practices** – This strategy discusses best practices for large migrations and provides use cases from customers across various industries.
- **Introduction to Database Migration** – In this course, you learn how to migrate a production database by using the AWS Database Migration Service (AWS DMS) and AWS Schema Conversion Tool (AWS SCT).
- **AWS DataSync Primer** – The course helps you get started with DataSync, showing you how to move large amounts of data between on-premises storage and the AWS Cloud.
- **Lift-and-Shift Application Workloads** – This webpage helps you understand the basics the rehost, or lift-and-shift, migration strategy.
- **AWS Application Migration Service (AWS MGN) – A Technical Introduction** – This course introduces the Application Migration Service.

- [Portfolio discovery and analysis for migration](#) – This guide defines the approach for defining, collecting, and analyzing the data required to create a migration plan.
- [Application portfolio assessment strategy for AWS Cloud migration](#) – This AWS Prescriptive Guidance strategy helps you understand the key stages to successfully assess your application portfolio.
- [AWS Cloud Migration Factory Solution](#) – This webpage helps you understand what AWS Cloud Migration Factory Solution is.
- [CloudEndure Migration Factory best practices](#) (YouTube video) – This video provides an overview of the AWS Cloud Migration Factory Solution and shares best practices for large-scale migrations. It includes information about how to coordinate and automate many manual migration processes.

Advanced training

Advanced training for large migrations dives deeper into the migration methodologies, tools, and best practices by providing workshops and training resources for the workstreams. The following are recommended for this training level:

- [Cloud migration factory workshop](#) – This technical workshop provides information about how to accelerate a large migration by using automation and the migration factory model.
- [Guide for AWS large migrations](#) – This guide contains high-level information about performing a large migration and introduces the large migration playbooks.
- [Foundation playbook for AWS large migrations](#) (this guide) – Use this playbook to train workstreams about preparing the platform foundation and people foundation for a large migration.
- [Project governance playbook for AWS large migrations](#) – This playbook provides step-by-step instructions for setting up the project governance framework and providing continuous governance throughout the migration.
- [Portfolio playbook for AWS large migrations](#) – This playbook provides step-by-step instructions to help you build your application prioritization runbook, metadata management runbook, and wave planning runbook.
- [Migration playbook for AWS large migrations](#) – This playbook provides step-by-step instructions for preparing migration runbooks for each migration pattern and preparing migration task lists.

Create your training dashboard

In the [foundation playbook templates](#), you can use the *Dashboard template for training* (Microsoft Excel format) as a starting point for building your own training plan and tracking tool. You use a training plan to assign a training level to each workstream. You then use a training tracking tool to record each individual's progress toward completing the required trainings in their workstream.

1. On the **Prerequisites** spreadsheet, **Fundamentals** spreadsheet, and **Advanced** spreadsheet, add or remove workstreams as appropriate for your large migration project.
2. On the **Prerequisites** spreadsheet, update the training materials as needed for your use case. Define the appropriate training for infrastructure, networking, and data centers. We recommend working with your IT department to identify technology-level training that is appropriate for all personnel in your large migration project. This spreadsheet should contain the training materials that you want all members of every workstream to complete.
3. On the **Fundamentals** spreadsheet, update the training materials as needed for your use case, and identify which workstreams should train on each item listed.
4. On the **Advanced** spreadsheet, update the training materials as needed for your use case, and identify which workstreams should train on each item listed.
5. On the **Training** tracker spreadsheet, enter the name of each individual in your large migration project and their workstream.

6. As each individual completes the required training for their workstream, mark the training as **complete**.

Platform foundation

This section focuses on assessing the readiness of the on-premises infrastructure, preparing the AWS landing zone or reviewing the existing landing zone design, and identifying the migration tools needed. You review the common infrastructure, operations, and security questions that you should consider for building a platform. You document your answers and decisions as migration principles. As a result, you have a solid platform to achieve the scale and velocity required for large migrations.

This section includes the following topics:

- [Landing zone considerations for a large migration \(p. 24\)](#)
- [On-premises considerations for a large migration \(p. 31\)](#)
- [Document your migration principles \(p. 33\)](#)

Landing zone considerations for a large migration

A *landing zone* is a well-architected AWS environment that is scalable and secure. By establishing standards for the landing zone, such as defining the number of accounts and designing the subnets and security groups, you build a solid foundation. This foundation gives you the ability to enable, provision, and operate your environment for both business agility and governance at scale while accelerating your cloud adoption journey. For more information about landing zones and strategies for building them, see [Setting up a secure and scalable multi-account AWS environment](#).

In addition to the standard business, operational, security and compliance considerations for your landing zone strategy, you must consider how to facilitate a large migration. You must design the landing zone to support existing, on-premises workloads during the migration and after, in cases where some workloads remain on premises. This guide provides additional landing zone considerations that affect the migration velocity and overall migration timeline.

Typically, landing zones are designed and deployed to support new workloads in the AWS Cloud. This is because organizations are adopting AWS before making the decision to migrate a large number of existing applications. The benefit of this approach is that the organization gains valuable knowledge and skills in AWS before the large migration, but it can also lead to conflicts between the various stakeholders. Some stakeholders might want to modernize the application during the migration because they want to take advantage of cloud-native features. However, the common goal of a large migration is to achieve maximum migration velocity and ease the transition by migrating as many applications as possible without modifying the workload. You then modernize these applications after the migration is complete.

Some key factors of the landing zone that can affect your large migration program project are:

- Network bandwidth availability and management
- Account strategy for workload isolation and resource management
- Security and administrative controls for migrated workloads

This section reviews the infrastructure, operations, and security questions that you should consider when building your AWS landing zone. It also contains recommendations for how to design and deploy your landing zone to support a large migration project. As you answer the questions in this section, these

decisions become migration principles, which you document according to the instructions in [Document your decisions as large migration principles](#) (p. 33).

Infrastructure considerations

Have you considered?	Description	Actions
<p>How much data will you migrate per day and per week?</p>	<p>The desired migration velocity dictates the type of network connection and network throughput requirements. It also can affect the wave planning selection criteria.</p>	<p>After you have completed the portfolio assessment, determine the total amount of storage needed for all migrated resources in the cloud. Use this value to calculate the amount of time required to migrate the data using the current network bandwidth. You might need to increase the bandwidth to meet the migration timeframes, or you might need to use alternatives, such as AWS Snow Family solutions. In the foundation playbook templates, you can use the <i>Data replication calculator</i> (Microsoft Excel format) to calculate the required bandwidth for each migration wave.</p>
<p>What is the average write speed of the source servers in each wave?</p>	<p>The bandwidth required to transfer the replicated data is based on the write speed of the participating source servers. The amount of bandwidth required for server replication is the average write speed of your source servers multiplied by the number of servers in the largest wave.</p>	<p>During portfolio assessment, you need to determine the average number of data writes performed per by each server. In the foundation playbook templates, you can use the <i>Data replication calculator</i> (Microsoft Excel format) to understand the bandwidth required for migration traffic. The bandwidth required for migration traffic is in addition to the bandwidth used for normal business activity. After the migration is complete, you no longer need the additional bandwidth to support the migration activities.</p>
<p>Could additional network activities or existing infrastructure limit or reduce the replication speed?</p>	<p>If the network bandwidth also supports other business functions, these activities can reduce the amount of bandwidth available for replicating servers during the migration.</p>	<p>Early in the project lifecycle, carefully assesses and calculate the network bandwidth required to support all business activities. Consider the bandwidth needed for normal business activities, server replication, and new migration-related activities, such as syncing on-premises file shares with data on AWS.</p>

AWS Prescriptive Guidance Foundation
playbook for AWS large migrations
Infrastructure considerations

Have you considered?	Description	Actions
		<p>Providers might have long lead times to increase the network capacity, and you might need to upgrade the existing on-premises infrastructure. Consider whether any additional upgrades would be required as a consequence of upgrading the network infrastructure. Assessing bandwidth requirements early in the project provides time to make any necessary changes.</p>
<p>Does your current AWS subnet strategy meet the IP addressing requirements for migrating the on-premises workloads?</p>	<p>The number of servers and workload isolation requirements dictates the subnet strategy for your landing zone.</p> <p>Large migrations might require larger subnets than you expect. In a large migration, you group workloads in subnets similar to their setup in the on-premises infrastructure. To simplify the migration, larger, flatter subnet designs are preferred initially, and then, during modernization, you redesign the subnets as needed.</p>	<p>When the portfolio assessment has enough information about the infrastructure inventory, assess the on-premises network structure and incorporate it into the landing zone design as early as possible.</p>
<p>How many servers do you plan to replicate and migrate in parallel?</p>	<p>The size of the largest migration wave affects the subnet requirements and AWS service quotas.</p>	<p>Review the high-level migration plan, and use that to design your subnet. For example, if you have a plan to migrate 200 servers into one subnet, the Classless Inter-Domain Routing (CIDR) range for that subnet should have enough IP addresses to support the target number of servers. Also, increase the AWS service quota for each target account as needed.</p>

AWS Prescriptive Guidance Foundation
playbook for AWS large migrations
Infrastructure considerations

Have you considered?	Description	Actions
Have you identified the security group strategies for your migration resources?	Security groups are used to manage the inbound and outbound traffic for AWS resources. It is important to design security groups early in order to avoid delaying the migration.	In your runbook for application prioritization, review the migration strategies, and then design the security groups based on the migration strategies. For example, if the migration strategy is to rehost most of the workloads, consider a temporary, generic security group that supports migration cutover instead of refactoring the network and applying application-specific security groups.
Are there load balancers in use?	Typically, when migrating servers in an environment with load balancers, you need to assess the configuration of the load balancer and then migrate the load balancer. Migration options for the load balancer include using Elastic Load Balancing (ELB) or a partner appliance-based solution.	Assessment of load balancers needs to start early in the discovery phase in order to account for any custom configurations. In most environments, load balancer configurations are fairly standard, but some might have complex logic that determines whether you can migrate to ELB or a partner appliance-based solution.
Do any servers need to retain their source IP address?	The safest and easiest way to migrate servers to the cloud is to allocate new IP addresses to the migrated instances. In some situations, you might need to keep the same IP address as the source server. For example, a legacy application might have a hardcoded IP address that no one knows how to change.	<p>Keeping source IP addresses affects how you form move groups when wave planning. The most common approach is to migrate a whole subnet to AWS in a single move group because this makes routing and switching straight-forward at the network level.</p> <p>The following are key actions for keeping IP addresses:</p> <ul style="list-style-type: none"> • Carefully assess cross subnet communications between servers. • Decide how you will switch routing of IP addresses for migrated servers. Common options include switching a whole subnet or deploying a network technology that manages static IP routing on a server-by-server basis.

Have you considered?	Description	Actions
How much latency is acceptable between the source and AWS?	It is common to start the migration with VPN links because they can be set up quickly and then transition to a direct connection established using AWS Direct Connect. VPN links generally have higher and more variable latency, which affects data throughput and, more importantly, application response times.	If you are using a high or variable latency connection type, review each application's requirements and plan the migration waves accordingly. Plan to put applications that require low latency connections in later waves, when alternative connection types are available.

Operations considerations

Have you considered?	Description	Actions
Have you identified an AWS account strategy for your landing zone?	AWS best practices for a well-architected environment recommend that you should separate your resources and workloads into multiple AWS accounts. You can think of AWS accounts as isolated resource containers: they offer workload categorization and can reduce the blast radius in the event of a disaster.	In your runbook for application prioritization, review your selected migration strategies and use them to determine your account strategy. For example, if you want to migrate as quickly as possible and rehost is the most common migration strategy, fewer accounts is easier to manage. However, if your migration strategies require modernizing applications and you need to separate business units for compliance reasons, you should include one or more accounts for each business unit in your account strategy.
Do you need to switch monitoring tools during the migration? If so, is this part of the migration process, or does it occur before or after the migration?	Monitoring tools are critical for cloud operations. Your existing tools might not work in the cloud because of compatibility or licensing reasons. As part of the design, you need to decide which monitoring tools to use for the workload in the AWS Cloud.	Select a monitoring tool before starting the migration. Make sure the migration team incorporates instructions for setting up monitoring in the migration patterns. We recommend building an automation script that replaces or reuses the monitoring tools, as needed.
Have you identified application owners, and are they aware of any changes that must be made to the application so that it functions properly in the cloud?	Large migration is a transformation rather than just an infrastructure project. Include application owners early to support the migration. For example, application owners validate the wave plan, create	Work with a project management office and Cloud Enablement Engine team to align with application team leaders and make sure that communication is clear across all application teams. For more information about

AWS Prescriptive Guidance Foundation
playbook for AWS large migrations
Operations considerations

Have you considered?	Description	Actions
	test plans, and participate in the cutover.	communication and project transparency, see the Project governance playbook for AWS large migrations .
Have you selected a backup and recovery solution, and does it work with migrated workloads?	Backup and recovery tools are critical for cloud operations. Your existing tools might not work in the cloud because of compatibility or licensing reasons. As part of the design, you need to decide which backup and recovery tools to use for the workload in the AWS Cloud.	Select backup and recovery tools before starting the migration. Make sure the migration team incorporates instructions for setting up backup and recovery in the migration patterns. We recommend building an automation script that replaces or reuses the backup and recovery tools, as needed.
Have you identified all shared services and deployed them in the landing zone?	<i>Shared services</i> are services that support multiple applications, such as email, Active Directory, or shared database environments. You typically need to deploy shared services in the cloud before the migration so that migrated applications perform as expected.	Schedule a deep dive with the infrastructure team and application team leaders before completing the landing zone design. Review and confirm the list of shared services that you must deploy in the cloud before starting the migration. The most common shared services are Active Directory, network devices, Domain Name System (DNS), and infrastructure software.
Have you reviewed AWS service quotas for your target AWS Region and account?	Every AWS service has a service quota. Some of these quotas can be increased. It is important to review quotas before cutover. If insufficient resources are available, the cutover might fail.	Review the migration plan. For any target account that requires an increased service quota, request an increase. For more information and instructions, see AWS service quotas .
Do you need to upgrade your AWS Support plan?	AWS Enterprise support plan offers 24/7 phone support and faster response times than other plans. Because the cutover window is usually very short, having access to an experienced engineer to help resolve cutover issues can be critical to the success of a large migration.	Contact your AWS account team to discuss different support options and select the appropriate support plan for your large migration project.

Have you considered?	Description	Actions
Have you notified your AWS technical account manager (TAM) about your large migration plan?	The AWS Enterprise On-Ramp support team assigns a pool of Technical Account Managers (TAMs) who coordinate access to proactive programs, preventative programs, and AWS subject matter experts. Your TAMs can schedule availability of support resources as needed.	Notify your AWS technical account manager of your upcoming large migration project and share your migration plan. Your TAMs will make sure AWS support resources are available when needed. For example, your TAMs can schedule a support engineer during cutover, and the engineer can help mitigate technical issues and streamline the cutover.

Security considerations

Have you considered?	Description	Actions
Have you identified AWS Identity and Access Management (IAM) roles and policies for access management?	Manage identity and access for all members of your large migration project. By attaching IAM roles to the migrated resources and defining access policies, you control who can access the migrated resources in the cloud.	Work with the migration team to identify the roles and responsibilities. Determine which roles can access which AWS account, and identify the level of access that each role has. Work with the security teams to validate that the IAM roles are correct for each target AWS resource.
Are there any compliance requirements for your workloads?	Workloads might have different compliance requirements, such as the Health Insurance Portability and Accountability Act (HIPAA) or payment card industry Data Security Standard (PCI DSS). You must identify these requirements before the migration and plan for how to meet them.	Work with compliance team and portfolio team to identify the compliance requirements for each application, and design your target AWS account accordingly. For example, you might need to migrate some workloads to AWS GovCloud (US) or to a specific AWS Region. We recommend that you document the compliance requirements for each application so that you can use this information later in the application prioritization and wave planning process.
Does your security team need to review and approve any tools or services that you plan to use during the migration?	A large migration project to the AWS Cloud uses many services, such as AWS Application Migration Service, AWS Database Migration Service (AWS DMS), AWS DataSync, and portfolio discovery tools	Work with the migration team to identify all of the tools, services, and applications that you expect to use in the migration. Work with the security team to review the company policies and

Have you considered?	Description	Actions
	(such as Flexera One). Some organizations require that all new tools and services are approved before use.	approve these tools accordingly before the migration starts.

On-premises considerations for a large migration

On-premises infrastructure that supports your business operations must also be prepared for the large migration. By preparing the current infrastructure, you can help reduce the impact of the large migration to the business operations and application users.

This section reviews the infrastructure, operations, and security questions that you should consider when preparing your on-premises infrastructure for the large migration. As you answer the questions in this section, these decisions become *migration principles*, which you document according to the instructions in [Document your decisions as large migration principles \(p. 33\)](#).

Infrastructure considerations

Have you considered?	Description	Actions
Have you designed the on-premises DNS and routers to support traffic to and from target AWS accounts?	Because of the large number of servers and target AWS accounts, it is important to confirm that different networking components are configured correctly to support the migration strategies and scale.	Review the design of routing tables, and make sure there are correct routes between the AWS accounts and on-premises data centers. Also, make sure the DNS server is able to support DNS queries from both on-premises servers and AWS resources.
How will the migration team access both the on-premises and AWS environments?	The migration team needs to access the source and target servers to perform migration activities, such as install a replication agent on a source server or uninstall old software on a target server.	Review the existing authentication and authorization mechanisms and build a strategy to grant access. You can use an Active Directory group, IAM role, and Security Assertion Markup Language 2.0 (SAML 2.0) federation to allow single sign-on to the AWS account. We recommend creating a local admin user in case there are any authentication issues with Active Directory.
Are there any known congestion points in the current network configuration that would slow data throughput during the migration?	A large migration requires lots of bandwidth to replicate the data from on-premises data center to the cloud. Understanding any existing congestion points or limitations helps you better plan the migration.	Review the network configuration with the networking team to better understand the network path from the source machines to the target AWS accounts. Identify potential congestion points, such as a connection that is shared between the migration and production workloads.

Operations considerations

Have you considered?	Description	Actions
Do you have any scheduled blocked days, also known as <i>change freezes</i> , that could impact the migration?	A change freeze during migration can take critical resources and time away from an ongoing migration project.	Review the change management process with the operations team, and take blocked days into consideration when you plan cutover windows.
Have you reserved change days for the migration?	Change management processes can be complex, and some organizations allow changes only in certain maintenance windows.	According to your change management process, schedule changes at least five waves in advance. This helps prevent delays.
Have all of the servers in scope for the migration been recently rebooted?	System changes or uninstalled patches might cause issues during the migration, which would necessitate long cutover windows or rolling back the server. The best practice is to confirm that the server has been recently rebooted on the target side before migrating.	Review the dates of the last server reboots. If a server has not been restarted within the last 90 days, schedule a restart before migrating the server.
How does the disaster recovery and business continuity plan work today, and has this been factored into the landing zone design?	Disaster recovery and business continuity plans are critical components of meeting the recovery time objective (RTO) and recovery point objective (RPO) of the application. You need to make sure these plans work for both your on-premises and AWS workloads during the transition period.	Review the existing disaster recovery and business continuity plans and make sure the plans work for your target AWS account. If not, design new plans before moving workload to the AWS Cloud.

Security considerations

Have you considered?	Description	Actions
Have you created firewall rules to support the large migration?	Depending on the processes in your organization, it can take a long time to complete a change request for firewall configurations.	Review the existing firewall change process with security team, and design a strategy for large migration firewall changes accordingly. You might need to design a custom process for the large migration project, or you might need to submit changes early in the project. It is recommended that you consider using an AWS virtual private cloud (VPC) as an extension to your data center and avoid

Have you considered?	Description	Actions
		building firewall rules that are too complex, which could significantly delay the large migration.
Have you set up Active Directory in the AWS environment?	Active Directory is used for authentication and authorization. You need to make sure the target account workloads are able to connect to the domain controller for authentication and authorization. You can either add a new domain controller in the target VPC, or you can allow the AWS workload to connect to the on-premises domain controllers.	Review the Active Directory design with your security and infrastructure teams. Make sure the target AWS account has connectivity to the correct domain controller. Make sure that the target AWS subnet CIDR blocks are in the correct Active Directory sites so that the workloads in AWS are able to connect to the nearest domain controllers.
Have you identified third-party connections and application interdependencies?	Third-party connections and application interdependencies require that you modify the firewall rule, network access control list, and security group.	During the deep dive session with the application owners, review the external dependencies for each application. Submit a request to modify firewall rules and the network access control list and change security groups accordingly, based on the third-party dependency requirements.
Does your on-premises environment have any additional security tools that control access and processes running on the systems, such as CyberArk?	You might need to assess and update these security tools in order to allow the migration tools to function in the AWS landing zone.	Review the access policy in your source environment. If a security tool is being used in the access policy, confirm that the tool functions in the AWS Cloud, and then make sure that the migration team has access to both the source and target environments. If any changes are required, add these steps into your migration runbooks.

Document your migration principles

After reviewing the landing zone and on-premises considerations, you should document your answers and decisions. These become the migration principles that guide the rest of the project.

Do the following:

1. In the [foundation playbook templates](#), open the *Migration principles template* (Microsoft Word format).
2. Review the infrastructure, operations, and security considerations in the [Landing zone considerations for a large migration \(p. 24\)](#) and [On-premises considerations for a large migration \(p. 31\)](#) sections of this guide, and discuss the questions with the recommended teams.

3. Document the infrastructure, operations, and security decisions in your migration principles document. For examples of how to record these decisions, see the following table.
4. As needed for your use case, add new categories, items, and principles. For example, you might want to record migration principles for portfolio assessment or project management decisions.

The following is an example of how you might record your decisions to some of the questions in this guide.

Category	Item	Principle
Infrastructure	DNS server	Use Amazon-provided DNS as the primary DNS server for all Amazon Elastic Compute Cloud (Amazon EC2) instances. Set up a conditional forwarder that forwards queries to an on-premises DNS server.
	Security groups	Use a temporary security group to permit all standard infrastructure traffic between the source and target environments.
	EC2 instance types	<p>If utilization data is available from a discovery tool, such as Flexera One or modelizeIT, use this information to help determine the target instance type.</p> <p>If utilization data is not available, size the target instance based on the provisioned central processing unit (CPU) and memory of the on-premises infrastructure.</p>
Operations	Clean up	Servers remain in the staging area until the migration phase is complete, at the end of the hypercare period.
	AWS Backup	By default, the tag applied to each instance is <code>backup = true</code> . If backups are not required, the migration teams should change the tag to <code>false</code> .
	Monitoring	Use Amazon CloudWatch for monitoring of EC2 instances. After cutover, remove the existing monitoring agent from the target EC2 instances.
Security	Active Directory	Build a domain controller in each VPC, and link the

AWS Prescriptive Guidance Foundation
playbook for AWS large migrations
Document migration principles

Category	Item	Principle
		subnet of that VPC to your Active Directory site. For more information, see Designing the Site Topology . This configures all clients to use the correct domain controller.
	Server access	Users must retrieve a password from CyberArk to connect to the source machines.
	AWS Management Console access	Users must use federated login to access the AWS Management Console.

Resources

AWS large migration

Strategy

[AWS large-migration strategy and best practices](#)

Guide

[Guide for AWS large migrations](#)

Playbooks

- [Portfolio playbook for AWS large migrations](#)
- [Project governance playbook for AWS large migrations](#)
- [Migration playbook for AWS large migrations](#)

Additional references

- [AWS service quotas](#)
- [Cloud Enablement Engine: A Practical Guide](#)
- [Overview of Data Transfer Costs for Common Architectures \(blog\)](#)
- [Setting up a secure and scalable multi-account AWS environment](#)

For training resources, see the following sections of this document:

- [Prerequisites \(p. 21\)](#)
- [Fundamentals \(p. 21\)](#)
- [Advanced \(p. 22\)](#)

AWS Prescriptive Guidance glossary

[AI and ML terms \(p. 37\)](#) | [Migration terms \(p. 38\)](#) | [Modernization terms \(p. 42\)](#)

AI and ML terms

The following are commonly used terms in artificial intelligence (AI) and machine learning (ML)-related strategies, guides, and patterns provided by AWS Prescriptive Guidance. To suggest entries, please use the **Provide feedback** link at the end of the glossary.

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?"
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.
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.
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.
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.
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.
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 .
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?"
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).
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.
target variable	The value that you are trying to predict in supervised ML. This is also referred to as an <i>outcome variable</i> . For example, in a manufacturing setting the target variable could be a product defect.
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.
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: <i>Epistemic uncertainty</i> is caused by limited, incomplete data, whereas <i>aleatoric uncertainty</i> is caused by the noise and randomness inherent in the data. For more information, see the Quantifying uncertainty in deep learning systems guide.

Migration terms

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

7 Rs	<p>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:</p> <ul style="list-style-type: none">• 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.

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 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](#).

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 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](#).

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

	<p>Conversion Tool (AWS SCT). It analyzes database schemas and code objects, application code, dependencies, and performance characteristics, and provides assessment reports.</p>
business continuity planning (BCP)	<p>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.</p>
Cloud Center of Excellence (CCoE)	<p>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.</p>
cloud stages of adoption	<p>The four phases that organizations typically go through when they migrate to the AWS Cloud:</p> <ul style="list-style-type: none">• 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 <p>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.</p>
configuration management database (CMDB)	<p>A database that contains information about a company's hardware and software products, configurations, and inter-dependencies. You typically use data from a CMDB in the portfolio discovery and analysis stage of migration.</p>
epic	<p>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.</p>
heterogeneous database migration	<p>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.</p>
homogeneous database migration	<p>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.</p>
idle application	<p>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.</p>
IT information library (ITIL)	<p>A set of best practices for delivering IT services and aligning these services with business requirements. ITIL provides the foundation for ITSM.</p>

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 .
large migration	A migration of 300 or more servers.
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 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 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 <i>migration factory</i> 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 CloudEndure 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 strategy	The approach used to migrate a workload to the AWS Cloud. For more information, see the 7 Rs (p. 38) entry in this glossary and see Mobilize your organization to accelerate large-scale migrations .
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).
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 .

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 <i>people acceleration</i> , because of the speed of change required in cloud adoption projects. For more information, see the OCM guide .
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.
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 .
responsible, accountable, consulted, informed (RACI) matrix	A matrix that defines and assigns roles and responsibilities in a project. For example, you can create a RACI to define security control ownership or to identify roles and responsibilities for specific tasks in a migration project.
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.
service-level agreement (SLA)	An agreement that clarifies what an IT team promises to deliver to their customers, such as service uptime and performance.
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.
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.
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.

Modernization terms

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

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.
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, <i>Domain-Driven Design: Tackling Complexity in the Heart of Software</i> (Boston: Addison-Wesley

	<p>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.</p>
microservice	<p>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.</p>
microservices architecture	<p>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.</p>
modernization	<p>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.</p>
modernization readiness assessment	<p>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.</p>
monolithic applications (monoliths)	<p>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.</p>
polyglot persistence	<p>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.</p>
split-and-seed model	<p>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.</p>
strangler fig pattern	<p>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</p>

two-pizza team

example of how to apply this pattern, see [Modernizing legacy Microsoft ASP.NET \(ASMX\) web services incrementally by using containers and Amazon API Gateway](#).

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. For more information, see the [Two-pizza team](#) section of the [Introduction to DevOps on AWS](#) whitepaper.

Contributors

The following individuals contributed to this document:

- Chris Baker, Senior Migration Consultant
- Dwayne Bordelon, Senior Cloud Application Architect
- Dev Kar, Senior Consultant
- Wally Lu, Principal Consultant

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](#).

update-history-change	update-history-description	update-history-date
Updated name of AWS solution (p. 46)	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 (p. 46)	—	February 28, 2022