aws

Project governance playbook for AWS large migrations

AWS Prescriptive Guidance

AWS Prescriptive Guidance: Project governance playbook for AWS large migrations

Table of Contents

Project governance playbook for AWS large migrations

Amazon Web Services ([contributors](#))

February 2022 ([document history](#))

> **(i) Note**
>
> The project teams, roles, and workstreams referenced in this guide are described in the [Foundation playbook for AWS large migrations](#). We recommend completing the foundation playbook in advance of starting the project governance tasks in this guide.

Effective project governance is critical to the success of a large migration to the AWS Cloud. *Project governance* defines the rules, boundaries, and plans for completing the migration. Common project governance tools include a communication plan, benefit-tracking office, escalation plan, and quality gates for migration and cutover. By completing this playbook, you create and customize the governance that defines how to run your migration project.

In the third phase of a large migration, *migrate and modernize*, you refine your project governance model and create many of the tools and templates that you use during the migration. You should complete the assess and mobilize phases prior to starting this process. For more information about the phases of a large migration, see [Phases of a large migration](#) in the *Guide for AWS large migrations*.

This playbook provides a step-by-step approach to quickly develop an effective governance model for a large migration project. It describes project governance for a large migration, which spans both stages of the migrate phase, initialization and implementation:

- In stage 1, *initialize*, you assess team readiness and stand up the governance model. You define the processes and tools that govern your large migration project. At the end of stage 1, you have project governance tools that are customized for your own use case.

- In stage 2, *implement*, you use the tools you created in the previous stage in order to adhere to your project governance plan.

Guidance for large migrations

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

The following figure shows the other documents in this series. Review the strategy first, then the guides, and then proceed to the playbooks. To access the complete series, see Large migrations to the AWS Cloud.
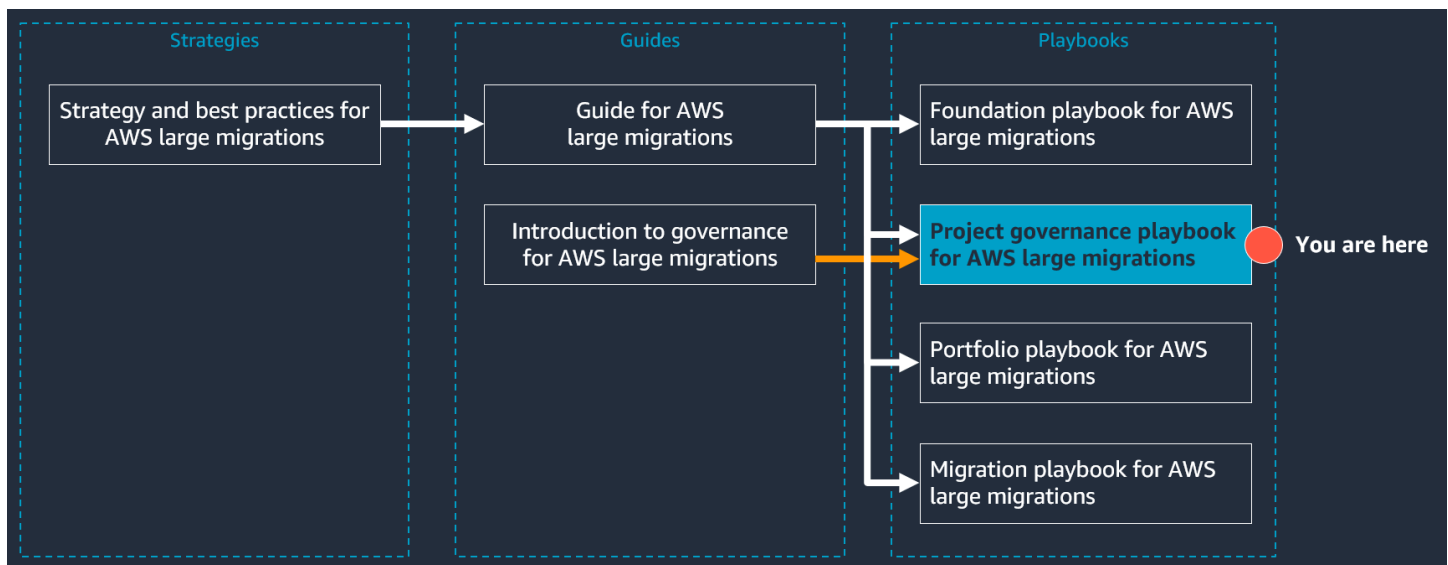
About the tools and templates

In this playbook, you create the following tools. You use these tools to communicate with the project stakeholders, including the migration teams, application owners, project sponsors, and executive leadership. The goal of the following tools is to maximize transparency for all project activities, which helps to accelerate the large migration:

- Kickoff presentation
- Meeting plan, including types and cadence
- Escalation plan
- Weekly project status report
- Wave workshop

- Cutover readiness assessment presentation

- Steering committee status report

- Benefit-tracking office

- Project summary dashboard

- Financial reporting process

- Resource plan

- Decision log

- Risks, actions, issues, and dependencies (RAID) log

- Communication plan and templates, such as gate communications and reminders

We recommend using the project governance playbook templates included in this playbook and then customizing them for your portfolio, processes, and environment. The templates are designed to foster effective communication, set clear expectations, and align executive leadership, application owners, and migration project stakeholders. The instructions in this playbook provide context as to the purpose of each of these templates, which your team can customize. This playbook includes the following templates:

- **Cutover readiness assessment template** – This template helps you track the progress of each wave through the quality gates and key project management milestones.

- **Financial glide path template** – This template is used to review financials with your project sponsors on a regular cadence.

- **Kickoff presentation template** – You use this presentation template at a kickoff meeting early in stage 1.

- **Meeting plan template** – You use this template to define the types of recurring meetings, establish their cadence, and identify the key participants.

- **Status report template** – You use this template in order to create a standard presentation format for the project status review meetings.

- **Steering committee meeting template** – You use this template in order to create a standard presentation format for the steering committee meetings.

- **Gate communication templates** – You use these email communication templates to share the status of the wave with project stakeholders and inform them of recent changes or upcoming activities. This playbook includes the following templates:

  - Communication template for cutover complete

- Communication template for hypercare complete
- Communication template for T-0
- Communication template for T-1
- Communication template for T-7
- Communication template for T-14
- Communication template for T-21
- Communication template for T-28

About managing a large migration

To manage and effectively govern a large migration project, the project manager needs to have a high-level understanding of the portfolio, the phases of a large migration, and the responsibilities of each workstream.

This section contains the following topics:

- Workstreams in a large migration
- Feeding the migration pipeline
- Hypercare period
- Establishing an agile approach

Workstreams in a large migration

In the migrate phase, at any given time, a minimum of four workstreams are operating simultaneously: the foundation, project governance, portfolio, and migration workstreams. These are the core workstreams of any large migration project, and your project might have additional, supporting workstreams. For more information, see Workstreams in a large migration in the *Foundation playbook for AWS large migrations*.

Feeding the migration pipeline

In the migration factory, wave planning and migration occur at the same time and operate continuously. The portfolio team feeds the migration pipeline by planning waves, and the migration team completes the pipeline by performing the migration and cutting over workloads. The portfolio team prepares five waves at the end of the initialization stage, and the implementation stage begins when the migration team begins migrating one or more of the prepared waves.

For each wave, the portfolio workstream runs 1–2 weeks, and the migration workstream typically runs 3–4 weeks. The portfolio workstream is five waves ahead of the migration workstream, so there is always a five-wave buffer between the portfolio and migration workstreams. Throughout the implementation stage, both the portfolio team and the migration team continue to process waves, and the buffer prevents the migration workstream from running out of servers to migrate. For an example of a wave schedule, see Stage 2: Implementing a large migration in the *Guide for AWS large migrations*.

The portfolio team prioritizes applications and then assigns them to waves in logical move groups. When planning waves, the portfolio team considers migration complexity, application similarities, and application and infrastructure dependencies. This helps make sure that the applications and their dependencies are migrated in their entirety. For more information about wave planning, see the [Portfolio playbook for AWS large migrations](#). For project governance, you manage and track information about the waves and sprints, including the applications, servers, and application owners. You might use a dashboard on a Confluence site, a list in Microsoft Excel, or a combination of tools.

Hypercare period

After you have completed the cutover, the migrated applications and servers enter the *hypercare period*. In the hypercare period, the migration team manages and monitors the migrated applications in the cloud in order to address any issues. Typically, this period is 1–4 days in length. At the end of the hypercare period, the migration team transfers responsibility for the applications to the cloud operations (Cloud Ops) team. At this time, the wave is considered complete.

Establishing an agile approach

By establishing an agile approach, the project team can remain flexible and quickly adapt to change during the migration. We recommend adopting a Scrum framework for a large migration. In the [Migration playbook for AWS large migrations](#), you assign waves to *sprints*, which is a fixed period of time in which the migration team works on all waves within that sprint. If each sprint is 2 weeks in duration, each wave spans at least two sprints. A sprint consists of standard events, such as planning the sprint and conducting daily stand-up meetings, a review, and a retrospective.

You use a *sprint backlog*, which consists of the current and pending tasks in the sprint, to manage the activities. In this playbook, you select a project management tool for tracking progress. You might select a project or issue-tracking application, such as Jira or Confluence, and you might also select a visual approach of representing tasks, such as a Kanban board or a Gantt chart. By tracking the sprint backlog in one or more of these tools, you provide project transparency, assign owners to each task, and establish clear deadlines.

Stage 1: Initializing a large migration

It is important to define the governance model early in the migrate phase and then conduct a kickoff meeting so that you can share it with the entire project team before you start migrating applications. If the governance model is already set up, skip to Stage 2: Implementing a large migration, where you will use the project governance tools and model established in stage 1. Establishing the right participants, communication formats, and meeting content at the outset allows you to focus on accelerating the migration. Ineffective planning for project meetings and communication can cause the team to spend too much time in meetings or providing status updates, rather than working on the migration.

> **ⓘ Note**
>
> The tasks in this chapter are intended to be performed concurrently. Many of the tasks are interdependent, as noted in the instructions for that task.

Stage 1 consists of the following sections, tasks, and steps:

- Before you begin
- Task: Kicking off the migrate phase
  - Step 1: Build a kickoff presentation
  - Step 2: Conduct the kickoff meeting
- Task: Creating a communications plan
  - Step 1: Create a communications team
  - Step 2: Establish an escalation plan
  - Step 3: Define meetings and their cadence
  - Step 4: Prepare meeting presentations
  - Step 5: Schedule recurring meetings for stage 1
  - Step 6: Understand the change management process
- Task: Defining communication gates and schedules
  - Step 1: Define the communication gates
  - Step 2: Create a T-minus schedule template
  - Step 3: Create standard email templates for each gate

- [Task: Defining project management processes and tools](#)
  - [Step 1: Select a project management tool](#)
  - [Step 2: Validate roles and responsibilities for all migration activities](#)
  - [Step 3: Establish a benefit-tracking office](#)
  - [Step 4: Create a project summary dashboard](#)
  - [Step 5: Create a financial reporting process](#)
  - [Step 6: Determine how to manage and scale resources](#)
  - [Step 7: Create a decision log](#)
  - [Step 8: Create a RAID log](#)

Before you begin

Confirm that you are ready to proceed with defining project governance for your large migration as follows:

- **Previous phases complete** – Defining project governance occurs in the third and final phase of a large migration. If you haven't already done so, we recommend you complete the assess and mobilize phases. For more information, see the [Guide for AWS large migrations](#).

- **Expertise available** – If you are new to a large migration project, have reviewed the available documentation, and would like support, consider engaging with in-house or external subject matter experts to prepare your team.

- **Migration team prepared** – Cutovers are likely to occur after regular working hours in order to minimize the impact on the business and users of the application. If this is the case with your project, confirm that the migration team and application owners are aware and prepared for the working schedule.

Task: Kicking off the migrate phase

To start the migrate phase of the project, you schedule a kickoff meeting. This meeting occurs once during the large migration project. Typically, you conduct this meeting as early as possible in stage 1, initializing a large migration. Aligning the project team members and setting expectations early helps the workstreams understand their responsibilities and build their runbooks. The purpose is to align the stakeholders and workstreams regarding the project scope, guiding principles, communication plan, and team member responsibilities.

In this task, you do the following:

- [Step 1: Build a kickoff presentation](#)
- [Step 2: Conduct the kickoff meeting](#)

Step 1: Build a kickoff presentation

In this step, you create a presentation for the kickoff meeting. As noted in the following steps, to create this presentation, you need some of the plans and processes that you define in other tasks in this playbook.

There are nuances to every project, but we recommend starting with the *Kickoff presentation template* (Microsoft PowerPoint format) available in the [project governance playbook templates](#). This template contains the core components, and you can customize it for your project. While you should review and customize the entire template, at a minimum, update the following slides:

1. On slide 4, define the project scope, guiding principles, factors that are critical to success, and the criteria by which success will be measured. You might work with a project management office, stakeholders, and the migration team to customize this slide for your organization.

2. On slide 5, create a roadmap of the high-level schedule for your project.

3. On slide 6, document the teams and key individuals involved in the migration. Identify individuals who are providing support from other teams in the organization, such as networking. Identify individuals by name and role, and differentiate internal and external resources. For a list of common roles in a large migration project, see [Roles](#) in the *Foundation playbook for AWS large migrations*.

4. On slide 10, add the T-minus schedules from [Step 2: Create a T-minus schedule template](#). Add new slides as needed to include a T-minus schedule for each migration strategy, such as replatform or refactor.

5. On slide 13, update the meeting plan according to [Step 3: Define meetings and their cadence](#).

6. On slide 16, add the escalation plan according to [Step 2: Establish an escalation plan](#).

7. On slide 20, add links to your shared repository and project management resources.

Step 2: Conduct the kickoff meeting

In this step, you schedule and conduct the kickoff meeting. Do the following:

1. Schedule the kickoff meeting to occur as early as possible in the migrate phase. Typical meeting participants include the project stakeholders, executive leadership, and the workstream leads.

2. Conduct the kickoff meeting and use the presentation you created in the previous step, Step 1: Build a kickoff presentation.

3. If there are any changes to the plans and processes presented in the meeting, after the meeting, update the plans accordingly.

4. Save the kickoff presentation in a shared repository so that all members of the large migration project can access the presentation as needed.

Task exit criteria

This task is complete when you have done the following:

- You have customized the *Kickoff presentation template* for your project.

- You have conducted the kickoff meeting.

- You have saved the kickoff presentation in a shared repository.

Task: Creating a communications plan

A critical element of the governance model is identifying who is responsible for communicating with application owners and how to escalate if an application owner doesn't respond. In this task, you define who is responsible for communications, determine what the regular communications and meetings will be, create your standard communication templates, and determine what happens if you need to escalate an issue.

In this task, you do the following:

- Step 1: Create a communications team

- Step 2: Establish an escalation plan

- Step 3: Define meetings and their cadence

- Step 4: Prepare meeting presentations

- Step 5: Schedule recurring meetings for stage 1

- Step 6: Understand the change management process

Step 1: Create a communications team

The communications team is part of the project governance workstream. This team is responsible for communicating with project stakeholders at key migration milestones, scheduling meetings, coordinating feedback, and confirming attendance for required meeting participants. The activities of the communications team are typically governed by communication gates, which you define in Task: Defining communication gates and schedules.

Do the following:

1. Identify the appropriate members of this team.

2. Designate a *communications lead*. This individual acts as a single point of contact throughout the migration for scheduling gate meetings, coordinating questions and feedback from the other workstreams, and confirming meeting attendance with required participants.

Step 2: Establish an escalation plan

When an issue arises in the migration, you must be able to quickly resolve it. By defining an escalation plan before the migration starts, you can provide a clear action plan to the team in advance, which helps prevent delay, frustration, or surprises. We recommend specifying a single-threaded leader for each business unit. If an application owner isn't engaging or responding, you can escalate to that individual.

This step is typically completed by the project manager and project sponsor. When establishing the escalation plan, you need to define the type of issue, the circumstances in which you should escalate the issue (known as the *trigger*), and define the tiers of escalation. We recommend no more than three tiers. For each tier, you should identify the *audience*, or *response owner*, and the amount of time that the audience has to respond. For example, if the first escalation audience doesn't resolve the issue within 24 hours, escalate the issue to the second tier, which is a different audience. With each escalation, CC the audiences of any prior tiers.

Do the following:

1. Create an escalation plan. You can use a dedicated project management tool for this, such as Jira or Confluence, or you can create a list in Microsoft Excel. We recommend documenting:

   - Short description of the anticipated or experienced issue

   - The trigger

- Tiers of escalation and audience

- The amount of time each tier has to respond to the issue

2. Conduct a meeting with the workstream leads and project sponsor to review the escalation plan.

3. Share the escalation plan with the entire project team to make sure that all members are familiar with the escalation process.

4. Save the escalation plan in a shared repository, and make sure that all project team members can access it.

#	Issue	Trigger	Tier 1		Tier 2		Tier 3
			Audience	*Escalate after*	*Audience*	*Escalate after*	*Audience*
1	Firewall ports need to be open to migrate workloads to AWS	Firewall isn't open by T-28 commit meeting	Network team, migration lead	24 hours	Network team manager	24 hours	Executive team, lead of impacted business unit

Step 3: Define meetings and their cadence

In this step, you identify the regular, recurring meetings for the migration project and establish the meeting frequency, or *cadence*. Documenting the meetings and their cadence improves project transparency. When an issue arises, team members can quickly identify the appropriate meeting to address it. You should identify the name of the meeting, the frequency, the core objectives, and the owners and participants. You might need to update this document as the migration progresses and you identify new meeting participants.

The following recurring meetings are common in a large migration project:

1. **Steering committee meetings** – These meetings are typically held twice a month, and the objective is to share the project status and resolve any issues that require involvement from

executive leadership. Participants of this meeting typically include the project sponsor, executive leadership, and a representative from the project management office.

2. **Project status review meetings** – These meetings are typically held once per week. The objective is to review the project status at the workstream level and evaluate the need for resources or subject matter experts. Participants of this meeting include the project manager, project stakeholders, workstream owners, and the migration lead.

3. **Daily stand-ups** – These are very short meetings held once per day. It is called a stand-up because the meeting should be short enough that the participants don't require a chair. The purpose is to review planned and recently completed tasks and surface any issues. In daily stand-ups, you typically use a visual task management tool, such as a Kanban board or Gantt chart, which you determine in Step 1: Select a project management tool.

4. **Infrastructure and operations checkpoint meetings** – These meetings are usually held twice a week. The objective is to review the progress of the migration, review active issues and decide whether escalation is required, collaborate across workstreams, and plan resources for the next sprint. Participants of this meeting include the technical team members who own RACI-defined migration activities.

5. **Migration business hours** – This time is reserved as an open meeting for application owners to seek support or guidance. We recommend that you hold business hours three times per week.

We recommend starting with the *Meeting plan template* (Microsoft Excel format) available in the project governance playbook templates. This template contains a default example, and you can customize it for your project.

Step 4: Prepare meeting presentations

As defined in Step 3: Define meetings and their cadence, large migrations require frequent meetings to align workstreams, address issues, and confirm that the migration is on schedule. Defining standard formats and presentations for these meetings helps participants by establishing consistent expectations for the meeting. It also helps reduce the amount of time needed to prepare for each meeting. In this step, you create the presentation templates for your regularly scheduled meetings.

We recommend starting with the following templates, which are included in the project governance playbook templates:

- *Status report template (Microsoft PowerPoint format)*
- *Steering committee meeting template* (Microsoft PowerPoint format)

- *Wave workshop template* (Microsoft PowerPoint format)
- *Cutover readiness assessment template* (Microsoft Excel format)

Do the following:

1. Customize the *Steering committee meeting template* for your project.

2. Customize the *Status report template* for your project. This presentation is used in project status review meetings, which are typically held on a weekly cadence. This template is a more robust version of the executive-level summary you created in the previous step.

3. Customize the *Wave workshop template* for your project. This presentation is used in the T-28 and T-14 commit meetings. In T-28 commit meetings, application owners commit to the wave, and in the T-14 commit meeting, they recommit to the cutover date.

4. Customize the *Cutover readiness assessment template* for your project. This presentation is used in the infrastructure and operations checkpoint meetings to review current progress of migration activities. The purpose of the presentation is to help the team confirm that the progress gates have been met and that the application is ready for cutover.

5. Store these presentation templates in a shared repository, where the meeting owners can access them.

6. For each type of meeting, define a shared repository where the meeting owners can save their presentations. After each meeting, the meeting owner should save a version of their presentation and any other meeting artifacts in this repository so that meeting attendees and the project team can reference this information. For example, the repository for the project status review meeting would contain a copy of the status report presented at each meeting.

Step 5: Schedule recurring meetings for stage 1

If you completed the mobilize phase, you might have already established some of the meetings in this step. Complete this step for any meetings that you haven't already scheduled. According to the meeting plan you developed in [Step 3: Define meetings and their cadence](#), the meeting owner should schedule the following recurring meetings:

- Daily stand-ups for each workstream
- Financial reporting meetings
- Steering committee meetings
- Project status reviews

- Infrastructure and operations checkpoint meetings

These meetings continue until the migration is complete.

Step 6: Understand the change management process

Understanding the change management process for your organization is critical to the success of a large migration project. The change management process affects the schedules and deadlines in your migration. You must understand the information and approvals required for each workload. Make sure that you understand:

- The deadlines for submitting the list of applications and servers in the wave plan

- The criteria and information required for gaining approval to move workloads on the planned date

- Any formal process documents that must be completed

- The process for submitting firewall or domain changes

All migration leads should understand the change management process prior to discovery activities. Some migration-related tasks require approval, and team members need to understand their responsibilities in the change management process. For more information about training, see Training and skills required for large migrations in the *Foundation playbook for AWS large migrations.*

Task exit criteria

This task is complete when you have done the following:

- You have created a communications team.

- You have defined participants for all meetings.

- You have established and approved an escalation plan.

- You have scheduled recurring meetings that start in stage 1, as defined in your meeting plan.

- You have defined the standard presentations that should be used in each meeting.

- For each meeting, you have defined a shared repository for capturing all presentations, activities, and artifacts.

- All change management processes are understood and documented.

Task: Defining communication gates and schedules

In stage 2 of a large migration project, the portfolio workstream is actively planning waves, and the migration workstream is migrating those waves. The project governance workstream oversees these activities and helps guide the waves through communication gates. A *communication gate* is a touchpoint when you formally communicate ongoing wave activities and status to the stakeholders. At each gate, a designated gate owner notifies the specified audience about the wave status and reminds application owners about upcoming activities or meetings. Gates typically correspond with migration milestones, and defining communication gates maximizes transparency for all project stakeholders. You move waves through the gates individually, or you can group waves together.

In this task, you do the following:

- [Step 1: Define the communication gates](#)
- [Step 2: Create a T-minus schedule template](#)
- [Step 3: Create standard email templates for each gate](#)

Step 1: Define the communication gates

During the migration, you repeat the communication gates for each wave or for a group of waves, until you have migrated all workloads and the project is complete. At a minimum, we recommend the following communication gates. You might decide to add more gates to your project as appropriate for your project.

Gate	Approximate timeline	Purpose	Gate owner	Audience
Gate 1: Create T-minus schedule	Before wave plan complete	Schedule dates for each gate	Project manager or communications team	Application owners, communication lead, migration lead
Gate 2: T-28 commit meeting	4 weeks prior to cutover	Kick off wave with application owners	Project manager or communications team	Application owners, communication

Gate	Approximate timeline	Purpose	Gate owner	Audience
				lead, migration lead
Gate 3: T-21 communication	3 weeks prior to cutover	Reminder that cutover is scheduled to occur in 21 days	Project manager or communica tions team	Applicati on owners, communication lead
Gate 4: T-14 checkpoint meeting	2 weeks prior to cutover	Review the schedule and assess progress on readiness tasks	Project manager and migration lead	Applicati on owners, communication lead, migration lead
Gate 5: T-7 communication	1 week prior to cutover	Reminder that cutover is scheduled to occur in 7 days	Communications team	Applicati on owners, operations team
Gate 6: T-1 go or no-go meeting	24–48 hours prior to cutover	Confirm readiness for migration cutover	Project manager or communica tions team	Cloud operation s team, applicati on owners, infrastructure team
Gate 7: T-0 cutover meeting	Day of cutover	Cut over and test the applicati ons	Project manager and migration lead	Cloud operation s team
Gate 8: Hypercare period start	1 business day after cutover	Notification that cutover is complete and hypercare period has started	Project manager or communica tions team	Application owners

Gate	Approximate timeline	Purpose	Gate owner	Audience
Gate 9: Hypercare period end	4 business days after cutover	Notification that hypercare period is complete	Project manager, communications team, or cloud operations team	Application owners in wave, communica tion lead, cloud operations team

The following image shows the sequence of these communication gates in the portfolio and migration workstreams. Gate 1 occurs during wave planning, gates 2–6 occur during the migration, gate 7 is the cutover meeting, and gates 8–9 occur during the hypercare period. Gates 2–6 are named with the format T-#. The T refers to time remaining, and the # is the number of days remaining until the scheduled cutover date.

Define the communication gates for your large migration project as follows:

1. Determine whether you require additional communication gates for your project. For example, if your project doesn't have a single-threaded leader who is responsible for facilitating migration readiness with application owners, you might want to include additional communication gates for reminding application owners about upcoming activities and due dates.

2. In a shared repository or project-tracking application, such as Jira or Confluence, record the communication gates for your large migration project. Make sure that you record the following attributes for each gate (for an example, see the communication gate table):

   - Gate number and name
   - Approximate timeline of when the gate occurs in relation to workstream milestones or cutover

- Purpose of the gate

- The individual or team who is responsible for the gate, known as the *gate owners*

- The individuals or teams who receive the communication or attend the gate meeting, known as the *audience*

- (Optional) The communication template or presentation template that the gate owner should use

Step 2: Create a T-minus schedule template

A *T-minus schedule* is a visual way to represent all of the high-level migration activities that need to be completed for each wave. It covers the period of time between the end of wave planning and the end of the hypercare period. Because the high-level migration activities vary based on the migration strategy, you need a T-minus schedule template for each migration strategy. You share the T-minus schedules at the kickoff meeting and at the T-28 and T-14 commit meetings.

Typically, you build a T-minus schedule by working back from the cutover date. You organize the activities into migration milestones, and you track detailed tasks separately within your project management tools. The T-minus schedule also highlights the communication gates that you defined in Step 1: Define the communication gates.

We recommend starting with the *T-minus schedule template* (Microsoft PowerPoint format), available in the project governance playbook templates. Do the following:

1. Open the *T-minus schedule template*. This template contains a default T-minus schedule for the rehost migration strategy.

2. Modify the default rehost migration activities based on your use case. For a list of activities for each migration strategy, refer to the responsible, accountable, consulted, informed (RACI) matrices that you created in the Foundation playbook for AWS large migrations.

3. Modify the default communication gates based on the decisions you made in Step 1: Define the communication gates.

4. Using the rehost T-minus schedule as a starting point, create a T-minus schedule for each migration strategy, such as replatform or refactor.

5. Share the T-minus schedules with the communications team, migration team, and cloud operations team. Make sure that all teams are in alignment and no adjustments are necessary.

6. Add the completed T-minus schedule templates to your kickoff presentation and to your wave workshop presentation.

Step 3: Create standard email templates for each gate

Create templates for the email communications that you will send to application owners at each communication gate. These emails should contain basic information about the applications in the wave, inform application owners of the wave status, and remind stakeholders about any upcoming due dates and meetings.

We recommend starting with the following templates, which are included in the [project governance playbook templates](#):

- *Communication template for T–28* (Microsoft Word format)
- *Communication template for T–21* (Microsoft Word format)
- *Communication template for T–14* (Microsoft Word format)
- *Communication template for T–7* (Microsoft Word format)
- *Communication template for T–1* (Microsoft Word format)
- *Communication template for T–0* (Microsoft Word format)
- *Communication template for cutover complete* (Microsoft Word format)
- *Communication template for hypercare complete* (Microsoft Word format)

Task exit criteria

This task is complete when you have done the following:

- You have defined the communication gates for your large migration project.
- You have created a T-minus schedule template.
- You have shared the T-minus schedule template with the project stakeholders.
- You have integrated the T-minus schedule template into your kickoff presentation and your wave workshop presentation.
- You have created standard templates for gate email communications.

Task: Defining project management processes and tools

Any large migration project requires well established management processes and tools. With a large migration, there are nuances to sharing information, tracking performance metrics, identifying the correct meeting participants, and assigning tasks to owners. In this task, you

document key migration tasks and owners, determine key performance indicators (KPIs) for the migration and decide how to measure them, track the budget, and develop tools for managing risks and tracking decisions.

Many of the steps in this task are performed concurrently, unless otherwise noted. Typically, you complete these steps before or just after the kick-off meeting.

In this task, you do the following:

- Step 1: Select a project management tool
- Step 2: Validate roles and responsibilities for all migration activities
- Step 3: Establish a benefit-tracking office
- Step 4: Create a project summary dashboard
- Step 5: Create a financial reporting process
- Step 6: Determine how to manage and scale resources
- Step 7: Create a decision log
- Step 8: Create a RAID log

Step 1: Select a project management tool

In this step, you establish the tools you want to use for tracking progress. You might choose to use a software solution such as Jira or Confluence, build your own dashboards in Microsoft Excel, or use a combination of these tools. Consider the following best practices when selecting or building project management tools:

- For tracking tasks and tracking progress, we recommend a visual management tool such as a Kanban board or Gantt chart, which are commonly available in project management applications. Visual management tools are particularly effective in the daily stand-up meetings for reviewing current tasks and wave progress.
- If you are selecting a project management application, consider whether you want to enter plans and processes (such as an escalation plan, decision log, or RAID log) in your project management tool, and make sure that it has the features you want.
- It is important that the project sponsor, executive leaders, project managers, and external stakeholders (if any) are aligned on the selected tool.

For more information about how these tools are used, see Establishing an agile approach.

Step 2: Validate roles and responsibilities for all migration activities

In the Foundation playbook for AWS large migrations, you created a detailed RACI matrix for each migration strategy and high-level task in your large migration project. A RACI matrix is a responsibility assignment tool, and the name is derived from the four responsibility types defined in the matrix: responsible (R), accountable (A), consulted (C), and informed (I). This matrix format is recommended to align roles and responsibilities across all of the migration activities. This matrix can align on-site teams with remote teams or external partners. In this step, you validate the matrices are correct and review them with the project teams.

In order to tailor the RACI tasks for your organization, we recommend you consider the following:

- Understand the change management processes, the lead times required for those processes, and the roles involved in approving changes. For more information, see Step 6: Understand the change management process.
- Make sure that you have vetted your back-up and disaster recovery strategy prior to starting the migration, and share this strategy with the migration team. If you identify gaps in the strategy, we recommend you use integrated cloud services, such as AWS Backup or CloudEndure Disaster Recovery.

Do the following:

1. If you haven't done so already, create a RACI matrix for each high-level task according to the instructions in the Foundation playbook for AWS large migrations.
2. Review the matrices with the respective teams in each matrix. Confirm that all detailed tasks are represented and that the teams are familiar with their responsibilities.
3. Update and create new matrices throughout the migration as you identify new migration strategies or supporting tasks.

Step 3: Establish a benefit-tracking office

This team is a small group of individuals who are responsible for assessing the migration against key performance indicators (KPIs). This team evaluates whether the migration is progressing according to schedule and can act on any delays or issues impeding progress. This team meets outside of the weekly or bi-weekly project status meetings.

In each meeting, this team usually reviews and answers the following questions:

- What is the current status of the migration?

- Are we on track to achieve our target outcomes?

- Are we measuring performance accurately?

- Do we need to make any adjustments in order to accelerate the migration?

If the benefit-tracking office determines that the migration isn't achieving the desired velocity, this team should recommend adjustments to process, resourcing, or communications plans.

Do the following to establish a benefit-tracking office for your large migration:

1. Identify the appropriate participants. Typical members of this team include the project sponsor, project manager, migration lead, and an empowered representative from each business unit that has workloads in scope.

2. Establish a regular meeting cadence for the benefit-tracking office. We recommend this team meets once every two weeks.

3. Define qualitative and quantitative KPIs for the large migration with the project sponsor, and collect input from executive leadership. The benefit-tracking office assesses the progress of the migration against your KPIs. Examples of KPIs include:

   - (Quantitative) Actual number of servers migrated compared to plan

   - (Quantitative) The number of servers decommissioned compared to plan

   - (Qualitative) Review of survey feedback and action plan

   - (Qualitative) Corrective steps made in response to survey feedback

Step 4: Create a project summary dashboard

The project team must work with key project stakeholders collectively to develop a dashboard that clearly conveys how the migration is progressing. Your project summary dashboard should do the following on a single page:

- Quantifies the overall completed and remaining workloads for the whole project

- Reflects the performance of the most recently completed wave (planned and actuals)

- Shows the anticipated workloads in the upcoming wave (planned)

We recommend starting with the *Project summary dashboard template* (Microsoft PowerPoint format), available in the [project governance playbook templates](). Do the following:

1. Modify the template as needed for your project. We recommend representing the allocation of servers to each migration strategy. The provided template includes the rehost and replatform migration strategies.

2. Review your project summary dashboard with the project stakeholders, including executive leadership, and make sure that all stakeholders are aligned and understand how to use and access the dashboard.

3. Save the dashboard in a shared repository. All stakeholders should be able access this information themselves as needed.

Step 5: Create a financial reporting process

Typically, you track financial reporting separately from the project status report because you want to provide it to a more limited audience. The financial report should include the *actual costs*, which are the costs incurred to date, and the *forecasted costs*, which are the expected costs for the remainder of the project. You track internal and external resource costs separately. To assess actual and forecasted internal resource costs, you can use in-house time reporting and your resource plan. For external resources, you should ask your partners or consultants to provide actual and forecasted costs.

We recommend starting with the *Financial glide path template* (Microsoft PowerPoint format), available in the [project governance playbook templates](#). Do the following:

1. Determine the stakeholders who should receive this financial report.

2. Determine whether this financial report will be shared in a meeting or through email.

3. Modify the template as needed for your project.

4. Review your financial report with the executive leadership team or project sponsors to confirm alignment on the format and content.

5. With the stakeholders, determine how frequently this report will be updated and reviewed.

6. Determine where you will save this financial report. Because it contains sensitive financial information, we don't recommend saving this template in the shared repository with the rest of the project documentation.

Step 6: Determine how to manage and scale resources

Managing resources effectively as the project progresses is critical to a large migration effort. As the project moves from the initialization stage into the implementation stage, the migration team must scale up in order to support the migration waves. At the same time, the discovery team might be able to start scaling down, depending on the remaining discovery activities. In this step, you map out the resource management and scaling plan for efficiency. This step is typically performed by the project manager and workstream leads. After the plan is defined, you audit constantly throughout the project to determine if you need all of the resources in the plan. For example, delays in building the migration pipeline or larger-than-anticipated waves would likely affect the resource plan.

The resource plan is different for each large migration, and it is typically determined by factors unique to your project. Common factors include project budget, how your project team is organized, how quickly discovery activities can be completed, how your portfolio is distributed to each migration strategy (such as refactor, rehost, or replatform), and how much time is needed for change management processes in your organization.

When planning resources, consider the migration strategies for your portfolio and how these affect your migration and portfolio teams. For example, rehost is a common strategy for large migrations because it is low complexity. Almost every large migration project has at least one rehost migration pod of 4–5 individuals. If you plan to include high-complexity migration strategies, such as replatform or refactor, you should create migration team pods for these strategies and include additional migration and portfolio team resources in your resource plan. For more information about workstreams, team structure, and how many people are needed for each pod, see Team organization and composition in the *Foundation playbook for AWS large migrations*.

In addition, the presence of specialized workloads, such as SAP, also warrants a separate, specialized team of individuals who have experience with those workloads. For more information about specialized workloads, see *MAP specialized workloads* at AWS Migration Acceleration Program.

Do the following:

1. Define the resources you need to support project governance. Typical resources include a program manager for delivery governance and oversight, a project manager, and a supporting project manager.
2. Define the resources you need to support migration tooling. Typical resources include a cloud architect or external consultant.

3. If your project includes migrating a specialized workload, such as an ERP system, define the resources you need to support that workload. Typical resources for a specialized workload include:

   - Project manager

   - Architecture lead

   - Architecture engineer

   - DevOps engineer

   - Specialized migration pod that contains:

     - Functional subject matter expert (SME)

     - Testing specialist

4. Define the resources you need to support each migration strategy, such as rehost. Typical resources include:

   - Project lead

   - Architects and engineers for compute, storage, and networking

   - Testing specialist

5. Allocate the number of resources needed to support these teams at various stages of the project, including discovery, initialization, and implementation. Consider the acceleration of the migration as you refine your processes, and consider how to scale resources down as you approach the end of a stage or project.

Step 7: Create a decision log

Throughout the large migration, leads make decisions to resolve any issues that arise. Because of the size and scope of a large migration project, the project manager cannot be present when every decision is made. Workstream leads are responsible for recording the decisions that affect their workstream. The project manager is responsible for reviewing the decisions and presenting recent decisions at the project status review meetings.

This step is typically performed by a project manager. In this step, you create a decision log in a shared repository and confirm that the workstream leads understand their responsibilities for logging decisions. When necessary, use the escalation plan to facilitate timely decision making. For more information, see Step 2: Establish an escalation plan. Confirm that all team members understand the types of decisions that can be made at each level.

Do the following:

1. Create a decision log. You can use a dedicated project management tool for this, such as Jira or Confluence, or you can create a list in Microsoft Excel. We recommend documenting:

   - Short description of the decision

   - Status

   - How the decision affects the project

   - Alternative options considered

   - Who made the decision

   - Date the decision was made

2. Conduct a meeting with the workstream leads to review the decision log and train them on how to use it. It is important that you establish a culture of recording decisions.

3. Save the decision log in a shared repository, and make sure that all workstream leads can access it.

4. Before each project status review meeting, review the log for any decisions made since the previous meeting, and include these decisions in your *project status report presentation*. This ensures project-level transparency for all decisions made over the course of the project.

Step 8: Create a RAID log

Similar to the decision log, you should track risks and issues in a project management tool known as a *risks, actions, issues, and dependencies (RAID) log*. No matter how thoroughly you plan your large migration, issues will occur, and you will identify some risks to your project. By identifying and recording risks and issues, you provide transparency to the project, and you establish a process to control and monitor potential issues, minimizing their impact to the project.

Do the following:

1. Create a RAID log. You can use a dedicated project management tool for this, such as Jira or Confluence, or you can create a list in Microsoft Excel. We recommend documenting:

   - Type (risk, action, issue, or dependency)

   - Short description of the item

   - Open date

   - Probability

   - Impact

   - Severity score, which is calculated by multiplying the probability and impact

- Owner

2. Conduct a meeting with the workstream leads to review the RAID log and train them on how to use it. It important that you establish a culture of recording risks and issues.

3. Save the RAID log in a shared repository and verify that all workstream leads can access it.

4. Before each project status review meeting, review the log for any risks and issues identified since the previous meeting, and include these in your *project status report presentation*. This ensures project-level transparency for all risks and issues.

Task exit criteria

This task is complete when you have done the following:

- You have selected one or more project management tools, such as Jira, Confluence, or dashboards and lists in Microsoft Excel.

- You have created and validated a detailed RACI matrix for each migration strategy (such as rehost) and for each high-level task in your large migration project.

- You have created a benefit-tracking office, established a regular cadence for their meetings, and created an ownership and reporting template for the meetings.

- Internal stakeholders are aligned for how financial reporting will be handled. You have established a formal cadence for reviewing the financial report, identified the recipients, and determined who should have access to the financial report.

- You have created a resource plan for your project.

- You have established a decision log in a shared repository, and all team leads are empowered to make updates.

- You have defined a location and template for the RAID log. You have established a process for maintaining the log and prioritizing issues. Week-to-week changes in the RAID log are summarized in the status report.

- All project stakeholders are aligned on how you will communicate the high-level project status in the project summary dashboard.

Stage 2: Implementing a large migration

In the previous stage, you established all of the tools, templates, plans, and processes that you need to govern the migration. In this stage, you use those assets to effectively manage and oversee the migration. This stage begins when the migration team begins migrating waves to the AWS Cloud. You repeat the gates in this stage for each wave or for a group of sequential waves.

Stage 2 consists of the following tasks:

- Task: Scheduling recurring meetings for stage 2
- Task: Completing the communication gates
  - Gate 1: Create a T-minus schedule for the wave
  - Gate 2: T-28 commit meeting
  - Gate 3: T-21 communication
  - Gate 4: T-14 checkpoint meeting
  - Gate 5: T-7 communication
  - Gate 6: T-1 go or no-go meeting
  - Gate 7: T-0 cutover meeting
  - Gate 8: Hypercare period start
  - Gate 9: Hypercare period end

Task: Scheduling recurring meetings for stage 2

According to the meeting plan you developed in Step 3: Define meetings and their cadence, the meeting owner should schedule the following recurring meetings. These meetings begin at the start of stage 2, after the first T-28 commit meeting, and continue until the migration is complete:

- Migration business hours
- Benefit-tracking office meetings

> ⚠ **Important**
>
> Continue to hold the recurring meetings that you set up in <u>Step 5: Schedule recurring meetings for stage 1</u>. These meetings continue until the end of the project.

Task: Completing the communication gates

In this task, you use the communication gates and the T-minus schedule that you defined in <u>Task: Defining communication gates and schedules</u> in order to communicate the status of each wave as it moves through the migration and portfolio workstreams.

You might move waves through these gates individually, or if multiple waves are on the same schedule, you can move them through the gates in a group. Because of the wave overlap in the migration workstream, at any given time in the migration, it is common to have multiple waves or groups of waves at different gates. The following table shows how waves overlap in the migration workstream, and each wave is scheduled 1 week apart. In this example, 6–7 waves are active in the migration workstream at any given time, and each wave is at a different gate.

Gate	Wave 1	Wave 2	Wave 3	Wave 4	Wave 5
Gate 1: T-minus schedule	March 13	March 20	March 27	April 3	April 10
Gate 2: T-28 meeting	March 20	March 27	April 3	April 10	April 17
Gate 3: T-21 communication	March 27	April 3	April 10	April 17	April 24
Gate 4: T-14 meeting	April 3	April 10	April 17	April 24	May 1
Gate 5: T-7 communication	April 10	April 17	April 24	May 1	May 8

Gate	Wave 1	Wave 2	Wave 3	Wave 4	Wave 5
Gate 6: T-1 go or no-go meeting	April 16	April 23	April 30	May 7	May 14
Gate 7: Cutover meeting	April 17	April 24	May 1	May 8	May 15
Gate 8: Hypercare period start	April 18	April 25	May 2	May 9	May 16
Gate 9: Hypercare period end	April 22	April 29	May 6	May 13	May 20

This task consists of the following communication gates:

- Gate 1: Create a T-minus schedule for the wave
- Gate 2: T-28 commit meeting
- Gate 3: T-21 communication
- Gate 4: T-14 checkpoint meeting
- Gate 5: T-7 communication
- Gate 6: T-1 go or no-go meeting
- Gate 7: T-0 cutover meeting
- Gate 8: Hypercare period start
- Gate 9: Hypercare period end

Gate 1: Create a T-minus schedule for the wave

Do the following in this communication gate:

1. Create a single, shared repository where you will store documentation for this wave.

2. Using the T-minus schedule template that you created in Step 2: Create a T-minus schedule template, enter dates specific to this wave, and then save the T-minus schedule in the shared repository.

3. Create a copy of the migration task list that you created in the Migration playbook for AWS large migrations, and then save it in the shared repository. You use this task list as a checklist as you progress through the gates.

4. Schedule the T-28 commit meeting with the appropriate participants. For more information about this meeting, see Step 3: Define meetings and their cadence.

Gate exit criteria

Continue to the next gate when you have completed the following project governance activities:

- You have established a shared repository for the wave.

- You have created a T-minus schedule for the wave.

- You have created a migration task list for the wave.

- You have scheduled the T-28 commit meeting.

Continue to the next gate when you have completed the following migration activities and any other tasks defined in your migration runbook:

- The portfolio team has completed the wave plan.

- The portfolio team has collected the migration metadata for the wave.

Gate 2: T-28 commit meeting

In this gate, the migration team reviews the wave plan with the application owners and asks the applications owners to commit to the wave plan and cutover date. Do the following in this communication gate:

1. Using the *wave workshop presentation* that you created in Step 4: Prepare meeting presentations, customize this presentation for the wave, and then save the presentation in the shared repository. You use this presentation in this gate and Gate 4: T-14 checkpoint meeting.

2. Conduct the T-28 commit meeting and, using your presentation, review the following:
   - Provide an overview of the wave plan and migration process.

- Provide detail about upcoming action items for the application owners.

- Confirm that the application owners are prepared to migrate each application in this wave.

- Confirm that application owners understand that they need to provide test plans for their applications. A *test plan* describes how to validate that the cutover was successful. Testing occurs immediately after cutover so that, if there are any issues, the migration team can roll back the application to its original environment with minimal impact to the business and application users.

- Review how stakeholders are expected to collaborate and communicate throughout the wave. Provide the location of the shared repository where stakeholders can find documents related to this wave.

- Review the escalation plan that you developed in Step 2: Establish an escalation plan.

- Provide an opportunity for questions and answers.

3. After the T-28 commit meeting, send the *T-28 communication email* that you created in Step 3: Create standard email templates for each gate. Customize the email for the wave information and recipients, and add all applications and servers in this wave.

4. After the T-28 commit meeting, schedule the following meetings with the appropriate participants:

   - T-14 checkpoint meeting

   - T-1 go or no-go meeting

   - T-0 cutover meeting

Gate exit criteria

Continue to the next gate when you have completed the following project governance activities:

- You have conducted the T-28 commit meeting.

- You have informed all key stakeholders about the shared repository to access wave documentation, and all stakeholders have access.

- You have started holding migration business hours, per Task: Scheduling recurring meetings for stage 2.

- Application owners have confirmed that the applications in the wave plan can be migrated.

- All stakeholders understand the communication approach and know which meetings they are required to attend.

- Application owners understand the specific action items for which they are responsible.

- You have sent the *T-28 communication email* to all stakeholders.

- You have saved the meeting presentation and meeting notes in the shared repository so that all stakeholders can access it.

- You have scheduled the T-14 commit meeting.

- You have scheduled the T-1 go or no-go meeting.

- You have scheduled the T-0 cutover meeting.

Continue to the next gate when you have completed the following migration activities and any other tasks defined in your migration runbook:

- You have updated the wave plan with any changes made during the T-28 commit meeting.

- You have submitted a request for change (RFC) for the applications and servers in the wave, and the change window is scheduled.

- Understand and identify the change management process.

- You have submitted RFCs for any new infrastructure requirements, such as forwarding, routing, or proxy services.

- You have updated the migration task list.

Gate 3: T-21 communication

The communications team continues to maintain contact with the application owners and business unit representatives. These stakeholders are invited to migration business hours to provide an opportunity for questions.

1. Send the *T-21 communication email* that you created in [Step 3: Create standard email templates for each gate](). Customize the email for the wave information and recipients, and add all applications and servers in this wave.

2. Update the scheduled T-14 checkpoint meeting with correct application owners. If any required participants cannot attend, confirm that an alternate representative can attend according to your escalation plan.

Gate exit criteria

Continue to the next gate when you have completed the following project governance activities:

- You have sent the *T-21 communication email* to all stakeholders.

Continue to the next gate when you have completed the following migration activities and any other tasks defined in your migration runbook:

- You have verified that the source servers meet the minimum requirements for replication.
- You have started replicating applications and servers in the wave.
- You have updated the migration task list.

Gate 4: T-14 checkpoint meeting

In this gate, you conduct the T-14 checkpoint meeting with the application owners and assess whether the team is on track to cut over as scheduled. Do the following in this communication gate:

1. Using the *wave workshop presentation* that you prepared in Gate 2: T-28 commit meeting, update the presentation for the T-14 checkpoint meeting.

2. Conduct the T-14 checkpoint meeting and review the following:

   - Review the applications and servers that are being migrated in this wave.

   - Review the remaining tasks and schedule to make sure that attendees understand the remaining steps in the process.

   - Confirm that all application owners (or their representative) are available for the cutover meeting.

   - Confirm that test plans are ready for when the cutover is complete.

3. After the T-14 checkpoint meeting, send the *T-14 communication email* that you created in Step 3: Create standard email templates for each gate. Customize the email for the wave information and recipients, and add all applications and servers in this wave.

4. Update the invitation to the T-1 go or no-go meeting and the T-0 cutover meeting with any changes in participants, such as an alternative representative designated by an application owner.

5. Update the migration task list.

Gate exit criteria

Continue to the next gate when you have completed the following project governance activities:

- You have conducted the T-14 checkpoint meeting. All application owners or their designated representatives attended. If an application owner didn't attend and is non-responsive, escalate the lack of attendance according to the escalation plan.

- You have conducted migration business hours for the week.

- You have sent the *T-14 communication email* to all stakeholders.

- You have saved the meeting presentation and meeting notes in the shared repository so that all stakeholders can access it.

- You have created a checklist of all pre-migration, migration, and post-migration tasks, closed any completed tasks, and saved the checklist in the shared repository.

Continue to the next gate when you have completed the following migration activities and any other tasks defined in your migration runbook:

- You have verified the health and status of replicated applications and servers. You are in the process of troubleshooting any issues or have completed troubleshooting.

- The application owners have provided test plans to the migration team.

- You have updated the migration task list.

Gate 5: T-7 communication

In this gate, the communications team continues to maintain contact with applications owners and business unit representatives. You also prepare for the cutover activities and meetings.

1. Send the *T-7 communication email* that you created in [Step 3: Create standard email templates for each gate](). Customize the email for the wave information and recipients, and add all applications and servers in this wave.

2. Confirm that required participants can attend the T-1 go or no-go meeting and the T-0 cutover meeting. Update the meeting invitations as needed to include alternate representatives.

Gate exit criteria

Continue to the next gate when you have completed the following project governance activities:

- You have sent the *T-7 communication email* to all stakeholders.

- You have confirmed attendance for the T-1 go or no-go meeting and T-0 cutover meeting. All participants have accepted the meetings, or alternative representatives have been identified.

Continue to the next gate when you have completed the following migration activities and any other tasks defined in your migration runbook:

- All requests for change for this wave have been approved.

- You have validated that the target infrastructure is ready for cutover.

- You have shut down any test instances that you created in order to validate the infrastructure.

- You have validated the cutover task list.

- You have updated the migration task list.

Gate 6: T-1 go or no-go meeting

In this gate, you review a checklist of pre-migration activities with all team members on the RACI matrix in order to validate that the applications and servers in the wave are ready for cutover. This gate occurs 24–48 hours before the scheduled cutover.

1. In the T-1 go or no-go meeting, review the checklist with all team members on the RACI matrix in order to validate that the applications and servers in the wave are ready for cutover.

2. Confirm that all required participants can attend the T-0 cutover meeting.

3. If you decide to proceed with migrating the wave (go), send the *T-1 communication email* that you created in [Step 3: Create standard email templates for each gate](). Customize the email for the wave information and recipients, and add all applications and servers in this wave.

4. If you decide not to proceed with migrating the wave or specific applications and servers (no-go), send an email to all stakeholders informing them of the decision and provide any available information about next steps or schedule changes.

Gate exit criteria

Continue to the next gate when you have completed the following project governance activities:

- You have confirmed that resources are available for the T-0 cutover meeting and that all required participants can attend.

- You have saved the meeting presentation and meeting notes in the shared repository so that all stakeholders can access it.

- You have sent the *T-1 communication email* to all stakeholders.

Continue to the next gate when you have completed the following migration activities and any other tasks defined in your migration runbook:

- In the migration task list, you have confirmed that all migration tasks are complete.

Gate 7: T-0 cutover meeting

In this gate, you migrate all servers and applications in the wave during a cutover meeting, and then you immediately have the application owners test the migrated applications to confirm that they are operating as expected. Application owners might attend the entire meeting or attend only as needed for their applications.

1. Before the cutover meeting, send the *T-0 communication email* that you created in Step 3: Create standard email templates for each gate. Customize the email for the wave information and recipients, and add all applications and servers in this wave.

2. In the T-0 cutover meeting, migrate the servers and applications in the wave according to the instructions in your migration runbooks, which you developed according to the instructions in the Migration playbook for AWS large migrations.

3. When an application or server has been migrated, use the test plan developed by the application owner to validate that the application is functioning as follows:

   - If the application or server is functioning as expected or has only minor issues, leave it in the AWS environment and remediate any issues.

   - If the application or server isn't functioning or has significant issues, roll it back.

4. As you complete the cutover activities in the migration task list, update the task list.

5. Send the *cutover complete communication email* that you created in Step 3: Create standard email templates for each gate. Customize the email for the wave information and recipients, and add all applications and servers in this wave.

Gate exit criteria

Continue to the next gate when you have completed the following project governance activities:

- You have validated that every application or server in the wave has been migrated successfully, or you have rolled it back.

- You have made note of any applications or servers that rolled back. For these applications or servers, you must update the migration pattern or redefine the target state to address any issues encountered during cutover. You will include these applications or servers in a future wave plan.

- You have sent the *cutover complete communication email* to all stakeholders.

Continue to the next gate when you have completed the following cutover activities:

- You have completed all steps in the *Cutover tasks* section of the migration task list.

Gate 8: Hypercare period start

In this gate, you do the following:

1. Ask project stakeholders to review the migrated applications and servers in the cloud. If any issues are identified, they should be sent to the migration team.

2. Address any issues identified during cutover or during the hypercare period.

3. Confirm that the cloud operations team is prepared to accept the workload.

4. Update all project management tools and repositories to reflect the status of the wave.

Gate exit criteria

Continue to the next gate when you have completed the following project governance activities:

- All stakeholders have reviewed the migrated applications and servers.

- The migration team has addressed any application or server issues that were identified during cutover or during the hypercare period.

- The cloud operations team has confirmed that they are ready to accept the migrated applications and servers.

- You have updated all project management tools and repositories in order to reflect the wave status.

Gate 9: Hypercare period end

The hypercare period typically lasts 1–4 days, and it ends when the migration team has resolved any issues with the migrated applications or servers. At the end of the hypercare period, the migration team meets with the cloud operations (Cloud Ops) team to review the migrated applications and servers. In this gate, the migration team transfers ongoing support of the migrated workloads to the Cloud Ops team. The Cloud Ops team notifies application owners that the hypercare period is complete and that they are now the point of contact for any issues. Optionally, you can include a survey in this communication and invite application owners to provide feedback about the migration and cutover process.

1. Incorporate the migrated applications and servers into the configuration management database (CMDB) for the cloud operations team.

2. Incorporate any application information into the Cloud Ops technical management support tool, such as ServiceNow.

3. Send the *hypercare complete communication email* that you created in [Step 3: Create standard email templates for each gate](#) for each gate. Customize the email for the wave information, and include instructions for how to contact the cloud operations team.

4. Notify the infrastructure support team of the transition in order to start the process of decommissioning the source servers and any supporting infrastructure. This step is typically performed by the Cloud Ops team or the project manager.

Gate exit criteria

This gate is complete when you have performed the following project governance activities:

- Cloud Ops has incorporated all workload-related information into their CMDB.

- Cloud Ops has incorporated all application information into their technical management support tool.

- You have sent the *hypercare complete communication email* to all stakeholders.

- The infrastructure team has started to decommission any supporting infrastructure that is no longer needed.

Resources

AWS large migrations

To access the complete AWS Prescriptive Guidance series for large migrations, see Large migrations to the AWS Cloud.

Additional references

- Mobilize phase (AWS Prescriptive Guidance)

Contributors

The following individuals contributed to this document:

- Pratik Chunawala, Principal Cloud Architect
- Bill David, Principal Customer Solutions Manager
- Wally Lu, Principal Consultant
- Amit Rudraraju, Senior Cloud Architect

Document history

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

Change	Description	Date
[Initial publication](#)	—	February 28, 2022

AWS Prescriptive Guidance glossary

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

Numbers

7 Rs

> Seven common migration strategies for moving applications to the cloud. These strategies build upon the 5 Rs that Gartner identified in 2011 and consist of the following:
>
> - Refactor/re-architect – Move an application and modify its architecture by taking full advantage of cloud-native features to improve agility, performance, and scalability. This typically involves porting the operating system and database. Example: Migrate your on-premises Oracle database to the Amazon Aurora PostgreSQL-Compatible Edition.
> - Replatform (lift and reshape) – Move an application to the cloud, and introduce some level of optimization to take advantage of cloud capabilities. Example: Migrate your on-premises Oracle database to Amazon Relational Database Service (Amazon RDS) for Oracle in the AWS Cloud.
> - Repurchase (drop and shop) – Switch to a different product, typically by moving from a traditional license to a SaaS model. Example: Migrate your customer relationship management (CRM) system to Salesforce.com.
> - Rehost (lift and shift) – Move an application to the cloud without making any changes to take advantage of cloud capabilities. Example: Migrate your on-premises Oracle database to Oracle on an EC2 instance in the AWS Cloud.
> - Relocate (hypervisor-level lift and shift) – Move infrastructure to the cloud without purchasing new hardware, rewriting applications, or modifying your existing operations. You migrate servers from an on-premises platform to a cloud service for the same platform. Example: Migrate a Microsoft Hyper-V application to AWS.
> - Retain (revisit) – Keep applications in your source environment. These might include applications that require major refactoring, and you want to postpone that work until a later time, and legacy applications that you want to retain, because there's no business justification for migrating them.

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

A

ABAC

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

abstracted services

    See [managed services](#).

ACID

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

active-active migration

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

active-passive migration

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

aggregate function

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

AI

    See [artificial intelligence](#).

AIOps

    See [artificial intelligence operations](#).

anonymization

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

anti-pattern

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

application control

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

application portfolio

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

artificial intelligence (AI)

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

artificial intelligence operations (AIOps)

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

asymmetric encryption

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

atomicity, consistency, isolation, durability (ACID)

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

attribute-based access control (ABAC)

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

authoritative data source

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

Availability Zone

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

AWS Cloud Adoption Framework (AWS CAF)

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

AWS Workload Qualification Framework (AWS WQF)

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

B

bad bot

A [bot](#) that is intended to disrupt or cause harm to individuals or organizations.

BCP

See [business continuity planning](#).

behavior graph

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

big-endian system

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

binary classification

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

bloom filter

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

blue/green deployment

A deployment strategy where you create two separate but identical environments. You run the current application version in one environment (blue) and the new application version in the other environment (green). This strategy helps you quickly roll back with minimal impact.

bot

A software application that runs automated tasks over the internet and simulates human activity or interaction. Some bots are useful or beneficial, such as web crawlers that index information on the internet. Some other bots, known as *bad bots*, are intended to disrupt or cause harm to individuals or organizations.

botnet

Networks of bots that are infected by malware and are under the control of a single party, known as a *bot herder* or *bot operator*. Botnets are the best-known mechanism to scale bots and their impact.

branch

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

break-glass access

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

brownfield strategy

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

buffer cache

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

business capability

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

business continuity planning (BCP)

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

C

CAF

> See [AWS Cloud Adoption Framework](#).

canary deployment

> The slow and incremental release of a version to end users. When you are confident, you deploy the new version and replace the current version in its entirety.

CCoE

> See [Cloud Center of Excellence](#).

CDC

> See [change data capture](#).

change data capture (CDC)

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

chaos engineering

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

CI/CD

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

classification

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

client-side encryption

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

Cloud Center of Excellence (CCoE)

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

cloud computing

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

cloud operating model

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

cloud stages of adoption

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

- Project – Running a few cloud-related projects for proof of concept and learning purposes

- Foundation – Making foundational investments to scale your cloud adoption (e.g., creating a landing zone, defining a CCoE, establishing an operations model)

- Migration – Migrating individual applications

- Re-invention – Optimizing products and services, and innovating in the cloud

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

CMDB

See configuration management database.

code repository

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

cold cache

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

cold data

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

computer vision (CV)

A field of AI that uses machine learning to analyze and extract information from visual formats such as digital images and videos. For example, Amazon SageMaker AI provides image processing algorithms for CV.

configuration drift

For a workload, a configuration change from the expected state. It might cause the workload to become noncompliant, and it's typically gradual and unintentional.

configuration management database (CMDB)

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

conformance pack

A collection of AWS Config rules and remediation actions that you can assemble to customize your compliance and security checks. You can deploy a conformance pack as a single entity in an AWS account and Region, or across an organization, by using a YAML template. For more information, see Conformance packs in the AWS Config documentation.

continuous integration and continuous delivery (CI/CD)

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

CV

> See [computer vision](#).

D

data at rest

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

data classification

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

data drift

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

data in transit

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

data mesh

> An architectural framework that provides distributed, decentralized data ownership with
> centralized management and governance.

data minimization

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

data perimeter

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

data preprocessing

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

data provenance

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

data subject

An individual whose data is being collected and processed.

data warehouse

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

database definition language (DDL)

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

database manipulation language (DML)

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

DDL

See [database definition language](#).

deep ensemble

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

deep learning

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

defense-in-depth

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

delegated administrator

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

deployment

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

development environment

See environment.

detective control

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

development value stream mapping (DVSM)

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

digital twin

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

dimension table

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

disaster

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

disaster recovery (DR)

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

DML

See database manipulation language.

domain-driven design

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

DR

See disaster recovery.

drift detection

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

DVSM

See development value stream mapping.

E

EDA

>   See [exploratory data analysis](#).

EDI

>   See [electronic data interchange](#).

edge computing

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

electronic data interchange (EDI)

>   The automated exchange of business documents between organizations. For more information,
>   see [What is Electronic Data Interchange](#).

encryption

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

encryption key

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

endianness

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

endpoint

>   See [service endpoint](#).

endpoint service

>   A service that you can host in a virtual private cloud (VPC) to share with other users. You can
>   create an endpoint service with AWS PrivateLink and grant permissions to other AWS accounts
>   or to AWS Identity and Access Management (IAM) principals. These accounts or principals
>   can connect to your endpoint service privately by creating interface VPC endpoints. For more

information, see Create an endpoint service in the Amazon Virtual Private Cloud (Amazon VPC) documentation.

enterprise resource planning (ERP)

A system that automates and manages key business processes (such as accounting, MES, and project management) for an enterprise.

envelope encryption

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

environment

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

- development environment – An instance of a running application that is available only to the core team responsible for maintaining the application. Development environments are used to test changes before promoting them to upper environments. This type of environment is sometimes referred to as a *test environment*.

- lower environments – All development environments for an application, such as those used for initial builds and tests.

- production environment – An instance of a running application that end users can access. In a CI/CD pipeline, the production environment is the last deployment environment.

- upper environments – All environments that can be accessed by users other than the core development team. This can include a production environment, preproduction environments, and environments for user acceptance testing.

epic

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

ERP

See enterprise resource planning.

exploratory data analysis (EDA)

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

F

fact table

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

fail fast

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

fault isolation boundary

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

feature branch

See branch.

features

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

feature importance

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

feature transformation

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

few-shot prompting

> Providing an LLM with a small number of examples that demonstrate the task and desired output before asking it to perform a similar task. This technique is an application of in-context learning, where models learn from examples (*shots*) that are embedded in prompts. Few-shot prompting can be effective for tasks that require specific formatting, reasoning, or domain knowledge. See also zero-shot prompting.

FGAC

> See fine-grained access control.

fine-grained access control (FGAC)

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

flash-cut migration

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

FM

> See foundation model.

foundation model (FM)

> A large deep-learning neural network that has been training on massive datasets of generalized and unlabeled data. FMs are capable of performing a wide variety of general tasks, such as understanding language, generating text and images, and conversing in natural language. For more information, see What are Foundation Models.

G

generative AI

A subset of AI models that have been trained on large amounts of data and that can use a simple text prompt to create new content and artifacts, such as images, videos, text, and audio. For more information, see What is Generative AI.

geo blocking

See geographic restrictions.

geographic restrictions (geo blocking)

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

Gitflow workflow

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

golden image

A snapshot of a system or software that is used as a template to deploy new instances of that system or software. For example, in manufacturing, a golden image can be used to provision software on multiple devices and helps improve speed, scalability, and productivity in device manufacturing operations.

greenfield strategy

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

guardrail

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

Detective guardrails detect policy violations and compliance issues, and generate alerts for remediation. They are implemented by using AWS Config, AWS Security Hub, Amazon GuardDuty, AWS Trusted Advisor, Amazon Inspector, and custom AWS Lambda checks.

H

HA

See [high availability](#).

heterogeneous database migration

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

high availability (HA)

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

historian modernization

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

holdout data

A portion of historical, labeled data that is withheld from a dataset that is used to train a [machine learning](#) model. You can use holdout data to evaluate the model performance by comparing the model predictions against the holdout data.

homogeneous database migration

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

hot data

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

hotfix

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

hypercare period

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

I

IaC

See infrastructure as code.

identity-based policy

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

idle application

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

IIoT

See industrial Internet of Things.

immutable infrastructure

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

inbound (ingress) VPC

> In an AWS multi-account architecture, a VPC that accepts, inspects, and routes network connections from outside an application. The AWS Security Reference Architecture recommends setting up your Network account with inbound, outbound, and inspection VPCs to protect the two-way interface between your application and the broader internet.

incremental migration

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

Industry 4.0

> A term that was introduced by Klaus Schwab in 2016 to refer to the modernization of manufacturing processes through advances in connectivity, real-time data, automation, analytics, and AI/ML.

infrastructure

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

infrastructure as code (IaC)

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

industrial Internet of Things (IIoT)

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

inspection VPC

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

Internet of Things (IoT)

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

interpretability

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

IoT

See Internet of Things.

IT information library (ITIL)

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

IT service management (ITSM)

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

ITIL

See IT information library.

ITSM

See IT service management.

L

label-based access control (LBAC)

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

landing zone

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

large language model (LLM)

A deep learning AI model that is pretrained on a vast amount of data. An LLM can perform multiple tasks, such as answering questions, summarizing documents, translating text into other languages, and completing sentences. For more information, see What are LLMs.

large migration

A migration of 300 or more servers.

LBAC

See label-based access control.

least privilege

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

lift and shift

See 7 Rs.

little-endian system

A system that stores the least significant byte first. See also endianness.

LLM

See large language model.

lower environments

See environment.

M

machine learning (ML)

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

main branch

See branch.

malware

Software that is designed to compromise computer security or privacy. Malware might disrupt computer systems, leak sensitive information, or gain unauthorized access. Examples of malware include viruses, worms, ransomware, Trojan horses, spyware, and keyloggers.

managed services

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

manufacturing execution system (MES)

A software system for tracking, monitoring, documenting, and controlling production processes that convert raw materials to finished products on the shop floor.

MAP

See Migration Acceleration Program.

mechanism

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

member account

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

MES

See [manufacturing execution system](#).

Message Queuing Telemetry Transport (MQTT)

A lightweight, machine-to-machine (M2M) communication protocol, based on the [publish/subscribe](#) pattern, for resource-constrained [IoT](#) devices.

microservice

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

microservices architecture

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

Migration Acceleration Program (MAP)

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

migration at scale

The process of moving the majority of the application portfolio to the cloud in waves, with more applications moved at a faster rate in each wave. This phase uses the best practices and lessons learned from the earlier phases to implement a *migration factory* of teams, tools, and processes to streamline the migration of workloads through automation and agile delivery. This is the third phase of the [AWS migration strategy](#).

migration factory

Cross-functional teams that streamline the migration of workloads through automated, agile approaches. Migration factory teams typically include operations, business analysts and owners,

migration engineers, developers, and DevOps professionals working in sprints. Between 20 and 50 percent of an enterprise application portfolio consists of repeated patterns that can be optimized by a factory approach. For more information, see the [discussion of migration factories](#) and the [Cloud Migration Factory guide](#) in this content set.

migration metadata

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

migration pattern

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

Migration Portfolio Assessment (MPA)

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

Migration Readiness Assessment (MRA)

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

migration strategy

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

ML

See [machine learning](#).

modernization

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

modernization readiness assessment

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

monolithic applications (monoliths)

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

MPA

See [Migration Portfolio Assessment](#).

MQTT

See [Message Queuing Telemetry Transport](#).

multiclass classification

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

mutable infrastructure

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

O

OAC

See [origin access control](#).

OAI

See [origin access identity](#).

OCM

See [organizational change management](#).

offline migration

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

OI

See [operations integration](#).

OLA

See [operational-level agreement](#).

online migration

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

OPC-UA

See [Open Process Communications - Unified Architecture](#).

Open Process Communications - Unified Architecture (OPC-UA)

A machine-to-machine (M2M) communication protocol for industrial automation. OPC-UA provides an interoperability standard with data encryption, authentication, and authorization schemes.

operational-level agreement (OLA)

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

operational readiness review (ORR)

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

operational technology (OT)

Hardware and software systems that work with the physical environment to control industrial operations, equipment, and infrastructure. In manufacturing, the integration of OT and information technology (IT) systems is a key focus for Industry 4.0 transformations.

operations integration (OI)

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

organization trail

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

organizational change management (OCM)

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

origin access control (OAC)

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

origin access identity (OAI)

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

ORR

See [operational readiness review](#).

OT

See [operational technology](#).

outbound (egress) VPC

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

P

permissions boundary

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

personally identifiable information (PII)

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

PII

See [personally identifiable information](#).

playbook

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

PLC

See [programmable logic controller](#).

PLM

See product lifecycle management.

policy

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

polyglot persistence

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

portfolio assessment

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

predicate

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

predicate pushdown

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

preventative control

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

principal

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

privacy by design

A system engineering approach that takes privacy into account through the whole development process.

private hosted zones

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

proactive control

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

product lifecycle management (PLM)

The management of data and processes for a product throughout its entire lifecycle, from design, development, and launch, through growth and maturity, to decline and removal.

production environment

See [environment](#).

programmable logic controller (PLC)

In manufacturing, a highly reliable, adaptable computer that monitors machines and automates manufacturing processes.

prompt chaining

Using the output of one [LLM](#) prompt as the input for the next prompt to generate better responses. This technique is used to break down a complex task into subtasks, or to iteratively refine or expand a preliminary response. It helps improve the accuracy and relevance of a model's responses and allows for more granular, personalized results.

pseudonymization

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

publish/subscribe (pub/sub)

> A pattern that enables asynchronous communications among microservices to improve scalability and responsiveness. For example, in a microservices-based [MES](#), a microservice can publish event messages to a channel that other microservices can subscribe to. The system can add new microservices without changing the publishing service.

Q

query plan

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

query plan regression

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

R

RACI matrix

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

RAG

> See [Retrieval Augmented Generation](#).

ransomware

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

RASCI matrix

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

RCAC

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

read replica

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

re-architect

See 7 Rs.

recovery point objective (RPO)

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

recovery time objective (RTO)

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

refactor

See 7 Rs.

Region

A collection of AWS resources in a geographic area. Each AWS Region is isolated and independent of the others to provide fault tolerance, stability, and resilience. For more information, see Specify which AWS Regions your account can use.

regression

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

rehost

See 7 Rs.

release

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

relocate

See 7 Rs.

replatform

See 7 Rs.

repurchase

See [7 Rs](#).

resiliency

An application's ability to resist or recover from disruptions. [High availability](#) and [disaster recovery](#) are common considerations when planning for resiliency in the AWS Cloud. For more information, see [AWS Cloud Resilience](#).

resource-based policy

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

responsible, accountable, consulted, informed (RACI) matrix

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

responsive control

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

retain

See [7 Rs](#).

retire

See [7 Rs](#).

Retrieval Augmented Generation (RAG)

A [generative AI](#) technology in which an [LLM](#) references an authoritative data source that is outside of its training data sources before generating a response. For example, a RAG model might perform a semantic search of an organization's knowledge base or custom data. For more information, see [What is RAG](#).

rotation

The process of periodically updating a secret to make it more difficult for an attacker to access the credentials.

row and column access control (RCAC)

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

RPO

See recovery point objective.

RTO

See recovery time objective.

runbook

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

S

SAML 2.0

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

SCADA

See supervisory control and data acquisition.

SCP

See service control policy.

secret

In AWS Secrets Manager, confidential or restricted information, such as a password or user credentials, that you store in encrypted form. It consists of the secret value and its metadata.

The secret value can be binary, a single string, or multiple strings. For more information, see [What's in a Secrets Manager secret?](#) in the Secrets Manager documentation.

security by design

A system engineering approach that takes security into account through the whole development process.

security control

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

security hardening

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

security information and event management (SIEM) system

Tools and services that combine security information management (SIM) and security event management (SEM) systems. A SIEM system collects, monitors, and analyzes data from servers, networks, devices, and other sources to detect threats and security breaches, and to generate alerts.

security response automation

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

server-side encryption

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

service control policy (SCP)

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

service endpoint

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

service-level agreement (SLA)

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

service-level indicator (SLI)

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

service-level objective (SLO)

A target metric that represents the health of a service, as measured by a service-level indicator.

shared responsibility model

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

SIEM

See security information and event management system.

single point of failure (SPOF)

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

SLA

See service-level agreement.

SLI

See service-level indicator.

SLO

See service-level objective.

split-and-seed model

A pattern for scaling and accelerating modernization projects. As new features and product releases are defined, the core team splits up to create new product teams. This helps scale your

organization's capabilities and services, improves developer productivity, and supports rapid innovation. For more information, see Phased approach to modernizing applications in the AWS Cloud.

SPOF

See single point of failure.

star schema

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

strangler fig pattern

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

subnet

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

supervisory control and data acquisition (SCADA)

In manufacturing, a system that uses hardware and software to monitor physical assets and production operations.

symmetric encryption

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

synthetic testing

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

system prompt

A technique for providing context, instructions, or guidelines to an LLM to direct its behavior. System prompts help set context and establish rules for interactions with users.

T

tags

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

target variable

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

task list

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

test environment

See environment.

training

To provide data for your ML model to learn from. The training data must contain the correct
answer. The learning algorithm finds patterns in the training data that map the input data
attributes to the target (the answer that you want to predict). It outputs an ML model that
captures these patterns. You can then use the ML model to make predictions on new data for
which you don't know the target.

transit gateway

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

trunk-based workflow

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

trusted access

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

tuning

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

two-pizza team

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

U

uncertainty

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

undifferentiated tasks

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

upper environments

See environment.

V

vacuuming

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

version control

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

VPC peering

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

vulnerability

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

W

warm cache

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

warm data

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

window function

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

workload

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

workstream

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

WORM

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

WQF

See [AWS Workload Qualification Framework](#).

write once, read many (WORM)

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

Z

zero-day exploit

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

zero-day vulnerability

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

zero-shot prompting

Providing an [LLM](#) with instructions for performing a task but no examples (*shots*) that can help guide it. The LLM must use its pre-trained knowledge to handle the task. The effectiveness of zero-shot prompting depends on the complexity of the task and the quality of the prompt. See also [few-shot prompting](#).

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.