



AWS Organizational Change Acceleration (OCA) 6-Point Framework – 2. Align Leaders

AWS Prescriptive Guidance



AWS Prescriptive Guidance: AWS Organizational Change Acceleration (OCA) 6-Point Framework – 2. Align Leaders

Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

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

Table of Contents

Introduction	1
Intended audience	3
Targeted business outcomes	3
About the OCA 6-Point Framework guides	3
2.1 IT and business leader alignment	5
Overview	5
Best practices	5
Example questionnaire	7
FAQ	7
Additional steps	8
2.2 Stakeholder assessment	10
Overview	10
Best practices	10
FAQ	11
Additional steps	12
2.3 Change impact assessment	14
Overview	14
Best practices	14
FAQ	15
Additional steps	17
2.4 Organization readiness assessment	18
Overview	18
Best practices	18
FAQ	19
Additional steps	20
2.5 Business case for change	21
Overview	21
Best practices	22
Creating a shared need for change	22
Shaping the vision	24
Writing your press release of the future and related FAQ	25
Cascading the case for change	26
FAQ	27
Additional steps	28

Resources 29

 References 29

 Partners 29

Contributors 31

Document history 32

Glossary 33

 # 33

 A 34

 B 37

 C 39

 D 42

 E 46

 F 48

 G 50

 H 51

 I 53

 L 55

 M 56

 O 60

 P 63

 Q 66

 R 66

 S 69

 T 73

 U 74

 V 75

 W 75

 Z 76

AWS Organizational Change Acceleration (OCA) 6-Point Framework – 2. Align Leaders

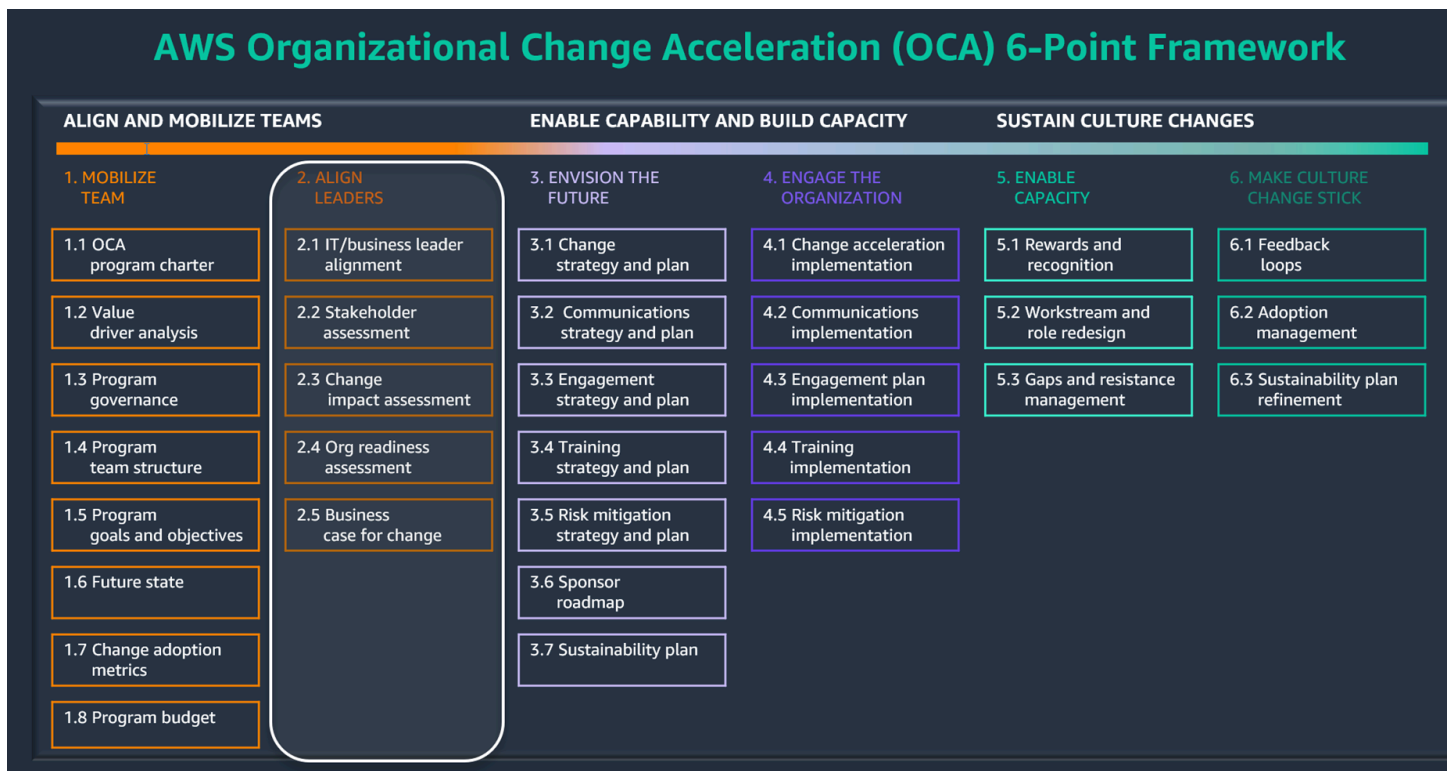
Amazon Web Services ([contributors](#))

January 2025 ([document history](#))

The AWS Organizational Change Acceleration (OCA) 6-Point Framework is intended to cover the full scope of people-related issues and challenges throughout the lifecycle of a cloud transformation, which might include migration, modernization, generative AI scaling, and innovation. This framework guides customer adoption of AWS technologies, processes, and new ways of working by:

- Identifying, aligning, and mobilizing key leaders
- Assessing and mitigating the organizational impacts of cloud transformation
- Designing change acceleration, communications, and training plans
- Developing leadership, sponsorship, and culture strategies

The framework's six points align with an agile sprint cadence, from program initiation through sustainable long-term change. The following diagram shows these six points and their subpoints.



Align Leaders is the second point. It helps you align and mobilize leaders around desired cloud outcomes, organizational impacts, and stakeholder readiness. *Align Leaders* contains five subpoints:

- [2.1 IT and business leader alignment](#). Establish shared understanding and commitment to cloud initiatives.
- [2.2 Stakeholder assessment](#). Identify impacted stakeholders, their influence, and disposition toward cloud adoption.
- [2.3 Change impact assessment](#). Analyze macro effects on skills, processes, and technology for each stakeholder group.
- [2.4 Organization readiness assessment](#). Evaluate the organization's ability to adapt to cloud transformation.
- [2.5 Business case for change](#). Create a compelling message that links cloud transformation to business rationale.

This guide discusses each subpoint of *Align Leaders* in detail.

Intended audience

This guide targets leaders who are responsible for accelerating cloud transformation. Following these recommendations will help minimize risks and maximize value.

Targeted business outcomes

The *Align Leaders* phase of the AWS OCA 6-Point Framework contributes to the following outcomes:

- **Value realization and return on investment (ROI):** Aligning IT and business leaders helps prioritize and align people-related factors to your cloud strategy and desired business outcomes.
- **Transformational leadership:** Leadership becomes aligned and mobilized to accelerate cloud transformation.
- **Cloud acceleration:** The alignment process establishes direction, metrics, governance, and program budget. These are all required to quickly mobilize resources for cloud transformation.
- **Organizational alignment:** The process works with leaders to establish desired business outcomes and specific goals to begin the change and start aligning organizational entities and performance levers.

About the OCA 6-Point Framework guides

This guide is part of a set of publications that cover the OCA 6-Point Framework, which is a programmatic and evidence-based organizational change adoption framework.

The content set includes a comprehensive set of templates, guidelines, supporting artifacts, assessments, accelerators, and tools that are designed to accelerate cloud transformation. We recommend that you start with the [overview](#) to understand the framework and its six points, and then consult the following individual guides for detailed discussions of each point.

1. [Mobilize Team](#)
2. Align Leaders (this guide)
3. [Envision the Future](#)
4. [Engage the Organization](#)
5. [Enable Capacity](#)

6. [Make Culture Change Stick](#)

For a comprehensive set of cloud transformation strategies, guidance, and resources, see [Accelerating cloud transformation](#).

2.1 IT and business leader alignment

Overview

Aligning IT and business leaders is crucial for successful cloud transformation. This alignment secures engagement, agreement, and funding from key stakeholders across global, regional, and functional areas. It builds sustained understanding and commitment to cloud initiatives, strategy, objectives, delivery plans, and change impact mitigation.

Key activities in IT and business leader alignment are:

- Stakeholder identification and interviewing
- Stakeholder management and alignment planning
- Leadership action planning
- Participation in key stakeholder updates

Best practices

Aligning IT with business leaders is critical to the success of your cloud transformation. Leaders will make decisions about the scope, budget, resourcing, and speed of the program. Their ability to align cohesively with IT will impact your internal and external customers.

Key best practices include:

- Onboard and prepare key stakeholders and leaders early.
- Identify areas of alignment and misalignment around strategic cloud objectives and change implications.
- Determine what leaders need to lead change effectively.

The alignment process identifies friction points and blockers to cloud adoption. Watch for organizational blockers such as:

- Misaligned priorities
- Resource constraints

- Budgetary concerns
- Leaders with outdated cloud knowledge
- Disengaged leaders
- Lingering impacts from mergers or acquisitions

Gather preliminary information before scheduling time to meet with leaders:

1. Review the business case for the cloud and any supporting data and documents such as a strategic plan, mission, vision, and press releases that might be available.
2. Review inputs such as a cloud strategy or roadmap, discovery information, Migration Readiness Assessment (MRA), and Migration Readiness Planning (MRP).
3. Identify key leaders to interview. Select stakeholders at levels that are high enough to have direct reports, a budget, and influence. Leaders should represent the global and functional footprint that is in scope for the cloud transformation.

At a minimum, involve the following individuals: executive sponsor, project leader, internal change team liaison, human resources (HR) lead, chief architect, data lead, security lead, operations lead, training lead, finance lead, infrastructure leaders, and lines of business leads.

4. Prepare a leadership alignment questionnaire. In general, this questionnaire should include about 7 to 10 open-ended questions that address perceptions of desired business outcomes, the relative priority of the cloud, sponsorship, and potential barriers.
5. Conduct leadership alignment interviews that are about 30 minutes in duration. At the beginning of the interview, establish the purpose of the interview and how the results will be used.
6. Analyze interview or survey data and develop a leadership assessment feedback report that shows areas of alignment and gaps.
7. Share the leadership assessment feedback report with the executive sponsor within one to two weeks of completing the analysis. Timeliness and bias for action are important to ensure that issues are quickly addressed and that the data stays relevant.
8. Work with the project sponsor to determine the next steps for closing alignment gaps among leadership team members.
9. Share the leadership assessment feedback report with the full leadership team and provide recommended next steps for building alignment.
10. Develop a plan for building IT and business leader alignment.

Example questionnaire

You can use the following questionnaire as an example for performing IT and business leader alignment interviews.

Start each interview by introducing yourself and the person who will be taking notes (if applicable), and ask the interviewee about their role, title, and years with the company. After these introductions, ask questions similar to the following:

- What is your sense of the reasons and rationale for your organization's cloud transformation? How well does your team understand those reasons?
- What outcomes do you expect?
- What changes of this magnitude have you experienced before? What was the outcome?
- How does your organization's culture encourage involvement in this kind of initiative?
- What effect will this cloud transformation have on your team's daily processes, roles, and responsibilities?
- What skills will need to change? What skills are missing?
- What barriers or risks do you perceive with this cloud transformation? Are there key blockers to overcome?
- Do you have any preferences or channels you recommend we use for communications and training?
- Who is the executive sponsor of this program? How will you personally sponsor this program within your team or organization?
- Do you have any other feedback you'd like to share?

FAQ

Q. What is IT and business leader alignment?

A. IT and business leader alignment is the process to identify, onboard, and prepare key stakeholders, target direct and indirect users of the cloud program, and mitigate the impacts associated with the journey to the cloud in a methodical fashion.

Q. Why is it valuable?

A. Aligning leaders is necessary to secure the engagement, agreement, and funding of key global, regional, local, and functional stakeholders to support and drive the cloud migration, modernization, and transformation efforts, and to transition to a new operating model. Aligning leaders builds sustained understanding and commitment to the initiatives and helps stakeholders understand the cloud strategy, objectives, delivery plan, and impacts.

Q. When do you conduct this activity?

A. To ensure successful cloud transformation, implement a robust IT and business leader alignment process within the first four to six weeks of program initiation. Conduct quarterly check-ins and reassess alignment after any significant organizational changes. Continuously monitor and address leadership gaps to maintain momentum and support throughout the transformation journey.

Q. Who is involved?

A. At a minimum, involve the executive sponsor, project leader, internal change team liaison, human resources (HR) lead, chief architect, data lead, security lead, operations lead, training lead, finance lead, infrastructure leaders, and lines of business leads.

Q. What are the inputs to this activity?

A. Inputs include a project charter, business case, cloud strategy, cloud readiness assessment results, and a list of key leaders (business and IT).

Q. What are the outputs of this activity?

A. The main output is an IT and business leadership assessment feedback report that summarizes the degree of alignment among leaders in terms of their understanding of cloud strategy, business case for change, priority of the cloud initiative, and support for cloud strategy. Additionally, the IT and business leader alignment activity might identify risks or potential blockers, leader perspectives on the business case for change, and specific leadership actions needed to advance cloud adoption.

Additional steps

To get started on IT and business leader alignment:

1. Identify leaders who are impacted and tangential to the program's success, timeline, resource planning, and budget.

2. Design a workshop for leaders to agree upon specific cloud transformation goals and the future state.
3. Determine a cadence at which these leaders will continually engage throughout the lifecycle of the program (for example, monthly, quarterly, during wave planning, during go/no-go decisions, for budget or scope approvals).
4. Begin to write and articulate a case for change based on the vision that the leaders have discussed, and use that message to create an introductory statement and communication campaign.
5. Determine if certain leaders need individualized touchpoints because of their influence on the program, and, if so, create leadership action plans and a cadence to review and make progress on those plans.
6. Evaluate the effectiveness of IT and business leader alignment periodically, and develop and implement leadership action plans as appropriate.

2.2 Stakeholder assessment

Overview

Stakeholder assessment is the first stage of managing stakeholders, to identify and understand their span of control, level of influence, and disposition toward cloud adoption.

A stakeholder assessment identifies and captures information about the people who will be impacted by the cloud program. This assessment can be used throughout the cloud migration or transformation journey to:

- Identify internal and external people who are affected by the change.
- Monitor readiness and potential challenges or risks.
- Support stakeholders throughout the cloud program.
- Identify change agents who will champion the cloud program.
- Understand the breadth and impact of the cloud program on the organization.

When you work with stakeholder groups, ask for guidance in segmenting and targeting audiences, preferred communication channels and key events, and contacts within the organization. You can use the insights you gained and the output of a stakeholder assessment to build communication plans, training plans, performance metrics, a network of change agents, and many more artifacts that last throughout the lifetime of the program. In addition, stakeholder assessment serves as a relationship-building opportunity and provides stakeholders with named contacts on the cloud team.

Best practices

The stakeholder assessment is reviewed regularly and updated throughout the cloud transformation to reflect changes in the project, its impacts, and the needs of its stakeholders. Stakeholders can be both organizations and people, but ultimately you must communicate with people. Be sure to identify the correct individual stakeholders within a stakeholder organization.

General considerations:

- Organizational characteristics and culture

- Regional compared with global segments
- Centralized compared with decentralized segments
- Language or translation requirements
- Other initiatives or events that are taking place or being planned for the key stakeholder group

The benefits of proper stakeholder assessment and management include:

- Early identification of powerful stakeholders
- Increased support and resources
- Better understanding of project benefits
- Anticipation of stakeholder reactions
- Early identification of conflicting objectives
- Increased employee and stakeholder engagement
- Targeted messaging and communications
- Improved communications and feedback
- Minimized change resistance

This assessment also helps the OCA team understand the following:

- Who will receive messages (the target audience)
- Who will help engage the target audience and deliver messages
- Who can ensure that messages translate into action
- Who to train and when, based on impact timing

FAQ

Q. What is the stakeholder assessment?

A. Stakeholder assessment is the first stage of managing stakeholders to identify and understand their span of control, level of influence, and disposition toward the cloud transformation effort.

Q. Why is it valuable?

A. It helps anticipate reactions, highlights perception gaps, and provides data to detect acceptance levels and attitudes toward the cloud program.

Q. When do you conduct this activity?

A. You should conduct the stakeholder assessment early in the program to inform the [business case for change](#), initial organizational readiness, and communication and training plans. You should review and update the assessment regularly throughout the cloud program to reflect changes in the project, scope, impacts, and stakeholder turnover (for example, people who leave or join the stakeholder group). On a routine basis, involve stakeholders in the ongoing management of the program.

Think of ways in which your team can involve stakeholders in program events, and ways in which stakeholders can involve the cloud program in their own events. The more employees are exposed to the cloud program through familiar communication channels from their own leadership, the more natural the transition to the cloud will be. As stakeholder engagement and interest in the cloud program increase, employees who report to each stakeholder will naturally engage, participate, and feel positive about the program.

Q. Who should be involved in the stakeholder assessment?

A. At a minimum, the executive sponsor, cloud leader, OCA leader, HR lead, chief architect, data lead, security lead, operations lead, training lead, finance lead, infrastructure leaders, and lines of business leads should be involved in the assessment.

Q. What are the inputs and outputs?

A. Inputs include the transformation vision, IT and business leader alignment assessment, and historical organizational data. Outputs include a report that provides a clear understanding of stakeholders' control levels, spheres of influence, and dispositions regarding cloud transformation.

Additional steps

To start the stakeholder assessment:

1. Review existing organizational information and cloud readiness assessments.
2. Prepare materials for the stakeholder assessment.
3. Identify and conduct stakeholder assessment with participants.
4. Identify key audience segments and their characteristics.

5. Develop a stakeholder assessment report.
6. Review findings with the cloud leadership team, executive sponsor, HR, and internal communications teams.
7. Use findings to formulate a strategy for communications and training.
8. Update the stakeholder assessment report regularly throughout the cloud adoption program.

2.3 Change impact assessment

Overview

A change impact assessment looks at the macro effects of the change and reports on the various skills, processes, performance management, and technology outcomes for each stakeholder group. This assessment is necessary to identify and capture significant differences between the current state and the desired future state. You can use this approach for any change effort to assess the magnitude of change.

Best practices

The change impact assessment should include:

- A gap analysis to understand and document the change (or gap) between the current and future states. For example, the gap might be a significant change in operational activities on premises compared with the cloud. In addition to identifying the changes, it is also important to document what stays the same.
- An assessment to understand the impact the change will have when it is implemented, based on the scale, scope, and size of impact (for example, the number of employees or business units that are affected).
- Documentation of potential areas of resistance (issues, risks, or barriers) that could prevent the change from being implemented successfully. This documentation helps you plan the activities in the change management plan and run them effectively. If there are a significant number of risks, you might have to document them in a separate change risk document.
- The identification of impacted stakeholder groups that will be the change targets or need to undergo a personal transition when the change happens.

The following questions facilitate the change impact identification process:

- How many people are impacted by the changes? Where are they located? What are their functions?
- How large is the gap between the current state and future state processes, tasks, and technologies used?
- Who will be impacted by this change (business units, functions, roles, locations, numbers)?

- Are there labor (union) issues associated with the change?
- How will impacted employees react to this change?
- What are the biggest barriers to implementing the change?
- What are the key impacts of the change (processes, technologies, people, and organizations)?
- What are the benefits of adopting the change?

The change impact assessment is typically documented in a template that's similar to the following:

Impacted area	Definition or description	Current state	Future environment	Change gap or impact	Who is impacted?	Level of impact	Change issues, risks, barriers
For example, leadership, culture, process, policies, structure, skills and capabilities, performance management, systems.	A brief overview of the change.	What's the current state?	What's the desired future state?	What are the major changes between the current and future states? What stays the same? What has to continue?	Who are the impacted stakeholders or change targets?	What is the level of change impact (for example, high, medium, or low).	What are the key issues or risks that could prevent the successful implementation of this change?

FAQ

Q. What is a change impact assessment?

A. It's an analysis of the macro effects of change on skills, processes, performance management, and technology for each stakeholder group.

Q. Why is it valuable?

A. It helps clarify changes at lower levels of granularity, determines appropriate steps for change acceleration plans, and identifies tangentially tied stakeholders.

Q. When should a change impact assessment be done?

A. It should be completed for any aspect of a cloud program where there's a substantial difference between current and future states for any stakeholder group. Here are some practical examples to consider:

- For managers, document when employees are likely to need training, when employees might need to have cloud-specific performance metrics incorporated into other annual performance plans, and when speaking points might be required.
- For HR stakeholders, document when key training events might be needed, when hiring plans might be required, how these changes might affect recruiting plans, when skill development opportunities become evident, when organizational design changes might be needed, and whether a compensation assessment should be conducted to market-test the value of cloud talent and skills.
- For working council or labor union stakeholders, document risks and concerns that might be raised and how best to address them, and if a regular meeting cadence should be established to improve transparency in communications.
- For finance stakeholders, document when a budget might be required for headcount and training activities, how budget processes and cycles might be affected by the cloud program, and how the transition from on premises to the cloud might change the way fixed and variable costs are treated in the company.

Q. Who should be involved in creating the change impact assessment?

A. Key participants should include the executive sponsor, cloud leader, OCA leader, HR lead, chief architect, data lead, security lead, operations lead, training lead, finance lead, infrastructure leaders, and lines of business leads.

Q. What are the typical inputs and outputs?

A. Inputs include the business case, process designs, organizational design models, readiness assessments, and subject matter expert (SME) interviews. Outputs include communication plans, training plans, stakeholder engagement plans, sponsor or leader plans, and updates to the business case, migration plan, and risk log.

Additional steps

To start the change impact assessment:

1. Define the process and tools.
2. Identify and document input sources.
3. Establish a cadence for capturing initial change impacts.
4. Brief leadership on findings and recommendations.
5. Update communication plans to address specific impacts and risks.
6. Engage HR if organizational restructuring or significant hiring needs are revealed.
7. Update training plans to address newly identified skill gaps.
8. Update overall change strategy to address identified impacts.

2.4 Organization readiness assessment

Overview

The primary objective in conducting an organization readiness assessment is to understand the organization's propensity, ability, and desire to adapt to change. It is also important to understand the organization's current culture and organizational structure as well as its desired state. This assessment helps identify strengths, barriers, and challenges in narrowing the gaps to achieving the future state.

Best practices

Before you deploy the assessment:

- Use existing employee pulse or culture surveys.
- Carefully consider the appropriate demographic data to gather.
- Choose the most suitable type of assessment for the organizational environment.
- Plan for follow-up assessments throughout the program to measure improvement.

The following table provides a list of sample questions that should be assessed on a 4-point Likert scale (strongly agree, agree, disagree, strongly disagree).

Pillar	Sample question
Leadership	Senior management (leadership levels above my team manager) actively supports the transformation.
Culture	For the transformation, mistakes will be treated as opportunities to learn instead of being punished as failures.
Training	I have learned the necessary skills to effectively work in the new environment.

Pillar	Sample question
Communications	A clear vision has been developed and communicated to employees regarding the transformation.

FAQ

Q. What is an organization readiness assessment?

A. It's a tool used to understand the organization's propensity, ability, and desire to adapt to change. This assessment is typically conducted through a survey.

Q. Why is it valuable?

A. It identifies opportunities and barriers, measures acceptance of change, and helps mitigate risks through action plans that support the overall objectives of the change effort.

Q. Who should be involved in this activity?

A. This activity must be conducted with the cloud leadership team, executive sponsors, and IT and business leaders.

Q. What are the inputs to this assessment?

A. Inputs include the business case, discovery phase outputs (MRA and MRP), interviews with the executive sponsor and HR team, staffing model, culture assessments, cloud strategy, and business value realization plans.

Q. What are the outputs of this assessment?

A. The primary output consists of baseline organizational readiness scores across the surveyed dimensions, along with prioritized mitigation plans that serve as inputs to the change strategy and plan.

Q. When do you conduct this assessment?

A. Conduct an organization readiness assessment at a key milestone such as a pilot application deployment. Use it periodically to measure progress and overall adoption of the change.

Q. How should you use the data from the assessment?

A. Use the findings to review the strategic vision and business case, obtain additional sponsorship, expand ownership to cross-functional leaders, invest in communications and training, and prioritize skill building.

Additional steps

To begin conducting an organization readiness assessment, follow these steps:

1. Review the strategic vision and business case.
2. Review historical survey data, if available.
3. Obtain sponsorship approval and support.
4. Determine the assessment tool and environment.
5. Review the question set with the executive sponsor for alignment, and determine anonymity.
6. Recruit the executive sponsor to send out the assessment.
7. Draft a letter from the executive sponsor that outlines the purpose and importance of the assessment.
8. Coordinate logistics (if necessary).
9. Conduct the assessment.
10. Compile and analyze results.
11. Develop a report that covers baseline scores and mitigation plans.

2.5 Business case for change

Overview

A business case for change is a compelling message that links cloud transformation to business rationale. It should:

- Be supported by a strong financial case.
- Consistently communicate the vision to generate stakeholder commitment.
- Be tailored for companywide or function-specific messages.
- Explain benefits to IT, business, finance, customers, and employees.
- Connect the cloud program to the external environment (for example, the competitive marketplace and customers).
- Establish a sense of urgency.

To test the case for change that you create, validate it against the following key criteria:

- The message communicates the future state in simple, clear terms. For example, "We are in the bottom quartile for introducing new products and losing x% market share. Our cloud transformation program will enable us to move to the top quartile, delighting shareholders and customers."
- The message describes why the change is necessary by describing the current state and explaining the consequences of starting or not starting the cloud transformation program at this time.
- The message describes how cloud transformation is aligned with other initiatives that improve business results.
- The message uses a metaphor to describe the future state, so it can be remembered easily.
- The message communicates your personal convictions (by tone or sentiment).
- The message describes what you will do personally to support the implementation of cloud transformation (for example, personal behavior changes, organizational systems changes).
- The message describes specific actions the audience can take to support the change (for example, attend cloud training sessions or workshops, establish a cloud leadership team, communicate to others about what is and isn't known about the cloud transformation program).

- The message is brief (5 minutes or less).

Best practices

- Create the case for change after conducting a stakeholder assessment.
- Articulate the benefits clearly and truthfully to influencers.
- Explain the consequences of not making the cloud journey.
- Use the case for change in various communications (for example, narratives, all-hands meetings).
- Tailor the message for specific audiences.
- Encourage employees to articulate the case for change to others.
- Engage in a two-way dialogue when you present the use case for change.
- Gather and respond to feedback and questions from all stakeholders.

Creating a shared need for change

Successful organizations learn to frame the need for change as more than a short-term threat. They find ways to communicate the need both as a threat and as an opportunity over the short term and long term. By doing so, they avoid communicating a short-term sense of urgency, and ensure the longer-term attention and involvement of key stakeholders. The key with motivation is to balance threats and opportunities to point and move people in the right direction.

Use the following table to collect the motivators for your business case for change. In the table:

- Threats ("if we don't change...") are reasons for change that focus on letting go of the current state. Threats present the current state as no longer attractive or even bearable.
- Opportunities ("if we do change...") are reasons for change that draw people toward a new or future state. Opportunities are future-focused and growth-oriented.
- Short-term motivators take effect relatively soon or quickly. The definition of short-term can be subjective and specific to the project or situation. Short-term motivators communicate a sense of urgency.
- Long-term motivators take effect sometime in the future or build up over time. They provide a sustainable power of motivation.

Short-term motivators:

Threats (if we don't change)	Opportunities (if we do change)
1.	1.
2.	2.
3.	3.
4.	4.
5.	5.
6.	6.
7.	7.

Long-term motivators:

Threats (if we don't change)	Opportunities (if we do change)
1.	1.
2.	2.
3.	3.
4.	4.
5.	5.
6.	6.
7.	7.

After you complete the tables, craft a 3–4 sentence statement of the need for change by using language that covers as many of the motivators in the table as possible.

Shaping the vision

An effective vision statement describes the outcome of the change. It is clear, legitimate, widely understood, and shared; the vision is shaped in behavioral terms. It is not a flashy slogan but describes what *good* looks like in the future. It is compelling, measurable, and emotionally exciting. It serves as the objective for all stakeholders who will be a part of, or impacted by, the change effort.

Use the following table to collect information for your vision statement.

Feedback from	What will you hear more of after cloud transformation?	What will you hear less of after cloud transformation?
Customers	1.	1.
Customers	2.	2.
Customers	3.	3.
Employees	1.	1.
Employees	2.	2.
Employees	3.	3.
Partners and suppliers	1.	1.
Partners and suppliers	2.	2.
Partners and suppliers	3.	3.

After you complete the table, craft a 3–4 sentence statement of the need for change by using language that covers as much of the feedback captured in the table as possible.

In addition, propose 3–5 success metrics that help you evaluate the achievement of the vision.

Writing your press release of the future and related FAQ

The press release of the future has its roots in innovation and new product development. The press release is written from the future point of view, when the new product will be released. An FAQ accompanies the press release and forces the writer to think broadly about the change. Using this approach has three key benefits: It helps you focus on the customer, it forces you to be explicit in your assumptions, and it can be interpreted by any stakeholder.

You can use this approach to ensure that your messaging is consistent, stakeholder-focused, and comprehensive.

Press release

Imagine that it is 12–18 months since you began your cloud transformation and you have been asked to speak at a press conference about the success of the cloud transformation journey and how it has solved customer needs, contributed to competitive positioning, enhanced employees' skills and careers, and added incremental revenue and income.

Use the following framework to write an article that might be written by the media after they hear your speech at the press conference.

If the press release is more than a page and a half, it is probably too long. Keep it short (three or four sentences for most paragraphs) and simple. You can include an FAQ with the press release to answer all other business or implementation questions, so you can keep the press release focused on customer benefits.

We recommend that you write your press release in the language of mainstream functional business leaders: Speak from the voice of your customer and avoid technical details.

The target audience for the press release are key stakeholders, who can be external customers or internal users of a solution, product, or service. The content centers on the customer problem, how current solutions (internal or external) fail, and how the cloud transformation will outperform existing solutions.

Here's an example outline for the press release:

- Title – Describe the most important benefits of the cloud transformation. (Keep it simple.)
- Subheading – Describe the benefits of cloud transformation for major stakeholder groups (for example, external customers, shareholders, and internal employees).

- **Summary** – Provide a summary of the business and financial outcomes of cloud transformation. Assume that the reader won't read further, so keep this section strong.
- **Problem or opportunity** – Describe problems or opportunities that are addressed by cloud transformation. (Copy information from the case for change.)
- **Solution** – Describe how cloud transformation addressed these problems or opportunities.
- **Company quote** – Provide a quote from a spokesperson in your company. (It could be you.)
- **Customer quote** – Provide a quote from a hypothetical customer that describes how they experienced the benefit.
- **Closing and call to action** – Wrap it up and provide links to additional resources.
- **FAQ** – Provide answers to questions that you anticipate. For example, here are some questions that relate to cloud transformation:
 - How will the customer experience change?
 - How will my role change?
 - How will the organization's culture change?
 - How will leadership responsibilities change?
 - What are the new skills required when we are in the cloud?
 - What new behaviors and mindsets are required?
 - What is the training plan?
 - What is the cloud transformation timeline?
 - What changes do we need to make to applications before the cloud transformation?
 - What communications are needed to share the future vision?
 - What communications are needed to inspire change adoption?
 - How will our support model change?
 - Do we need to change our development, test, integration, and delivery tools?
 - What level of server access will we have?
 - How will we monitor application performance?

Cascading the case for change

When you have the case for change articulated, think about how to communicate and roll it out to employees. You can use different communication formats for the case for change message. Here are some examples:

- A short slide deck that is shared in a large meeting such as an employee all-hands meeting or town hall
- An executive video where a key sponsor talks about the case for change and invites other leaders to talk about different aspects of the change holistically
- Posters or digital displays that are shared in company hallways, cafeterias, and break rooms
- Internal websites that talk about the program

FAQ

Q. What is a business case for change?

A. A case for change is a compelling, motivating, and inspiring message and document that ties the cloud transformation to the rationale for changing. Ideally, it is supported by a strong financial case and used to communicate the vision in a consistent way that generates commitment to cloud transformation from stakeholders. It can be tailored and expanded to communicate messages that are companywide or function-specific, and to explain the benefits to IT teams, business teams, finance teams, customers, and employees.

Q. Why is it valuable?

A. Leaders need to implement changes that will enable their organization to succeed in current and future markets. Employees might resist change if they don't believe in what leaders are asking them to achieve. There is a big difference in performance between employees who want to change and employees who change because they have to. A solid and well-communicated business case for change helps people commit to the cloud transformation journey voluntarily.

Q. When do you create it?

A. Create the business case for change early in the cloud program, and deliver it multiple times across all affected stakeholder groups.

Q. What are the inputs to this activity?

A. Inputs include project structure, goals, objectives, budget, metrics, stakeholder assessment, and change impact analysis.

Q. What are the outputs from this activity?

A. Outputs include key messages by audience, region, business unit, stakeholder group; change strategy and plan; communication strategy and plan; and training strategy and plan.

Q. Who is involved in this activity?

A. Participants include the executive sponsor, cloud leadership team, executive or steering committee, and leaders who participated in the [IT and business leader alignment](#).

Additional steps

To create the business case for change, follow these steps:

1. Review the case for change with others and get feedback.
2. Fine-tune the case for change based on feedback, and roll the plan out as necessary.
3. Assess the document's understandability, motivation, believability, and urgency.
4. Determine the appropriate audiences and sharing venues.

Resources

References

- [Accelerating your return on cloud investment by adopting a strategic transformation and change methodology](#)
- [AWS Change Acceleration 6-Point Framework and Organizational Change Management Toolkit](#)
- [AWS Organizational Change Acceleration \(OCA\) 6-Point Framework – 1. Mobilize Team](#)
- [AWS Organizational Change Acceleration \(OCA\) 6-Point Framework – 3. Envision the Future](#)
- [AWS Organizational Change Acceleration \(OCA\) 6-Point Framework – 4. Engage the Organization](#)
- [AWS Organizational Change Acceleration \(OCA\) 6-Point Framework – 5. Enable Capacity](#)
- [AWS Organizational Change Acceleration \(OCA\) 6-Point Framework – 6. Make Culture Change Stick](#)
- [AWS Cloud Adoption Framework \(CAF\)](#)
- [AWS Cloud Adoption Framework: People Perspective](#)

Partners

- **Accenture**
 - [Contact Partner](#)
 - [Contact the Accenture AWS Business Group](#)
 - [Future Talent Platform](#)
 - [Accenture and AWS take you further faster](#)
- **Deloitte**
 - [Contact Partner](#)
 - [AWS and Deloitte](#)
 - [Where Innovation Meets Impact](#)
- **PwC**
 - [Contact Partner](#)
 - [PwC and AWS](#)

- **Slalom**
 - [Contact Partner](#)
 - [AWS and Slalom Launch Centers](#)
- **Roberts Group Consulting**
 - [Contact Partner](#)

Contributors

- Melanie Gladwell, AWS Senior Practice Manager
- Scott Watson, AWS People Transformation Lead
- Tierra Jennings-Hill, AWS People Transformation Lead
- Nicole Lenz, AWS Sales Transformation Lead
- Jermel Moody, AWS Change Acceleration Lead

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	—	January 29, 2025

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

A2A (Agent-to-Agent)

A stateful protocol for agent-to-agent collaboration supporting task delegation and state transfer.

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 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.

Agent

An AI system that can autonomously reason, plan, and take actions using tools to achieve goals.

Agent Ops

Operational practices for building, testing, deploying, and running AI agents in production at scale.

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.

Citizen Developer

A business user who creates AI applications using no-code/low-code platforms without specialized technical skills.

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](#).

FM gateway

A centralized intermediary that controls and normalizes access to [foundation models](#). Also known as an *LLM gateway*.

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 CSPM, Amazon GuardDuty, AWS Trusted Advisor, Amazon Inspector, and custom AWS Lambda checks.

guardrails (AI)

Safety mechanisms that filter, validate, and constrain [agent](#) inputs and outputs to help ensure responsible and safe AI behavior.

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.

human-in-the-loop (HitL)

A workflow pattern where [agent](#) execution pauses for human review and approval at critical decision points.

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](#).

MCP

See [Model Context Protocol](#).

Model Context Protocol (MCP)

A stateless protocol for [agent](#)-to-[tool](#) communication.

MCP server

A service that exposes one or more [tools](#) through the [Model Context Protocol](#).

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.

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](#).

Shadow AI

Unauthorized [AI](#) applications built or used outside of governed channels within an organization.

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.

tool

A function or API that an [agent](#) can invoke to perform operations in external systems.

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.

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.