AWS Organizational Change Acceleration (OCA) 6-Point Framework – 1. Mobilize Team

AWS Prescriptive Guidance

AWS Prescriptive Guidance: AWS Organizational Change Acceleration (OCA) 6-Point Framework – 1. Mobilize Team

Table of Contents

AWS Organizational Change Acceleration (OCA) 6-Point Framework – 1. Mobilize Team

Amazon Web Services (contributors)

September 2024 (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.

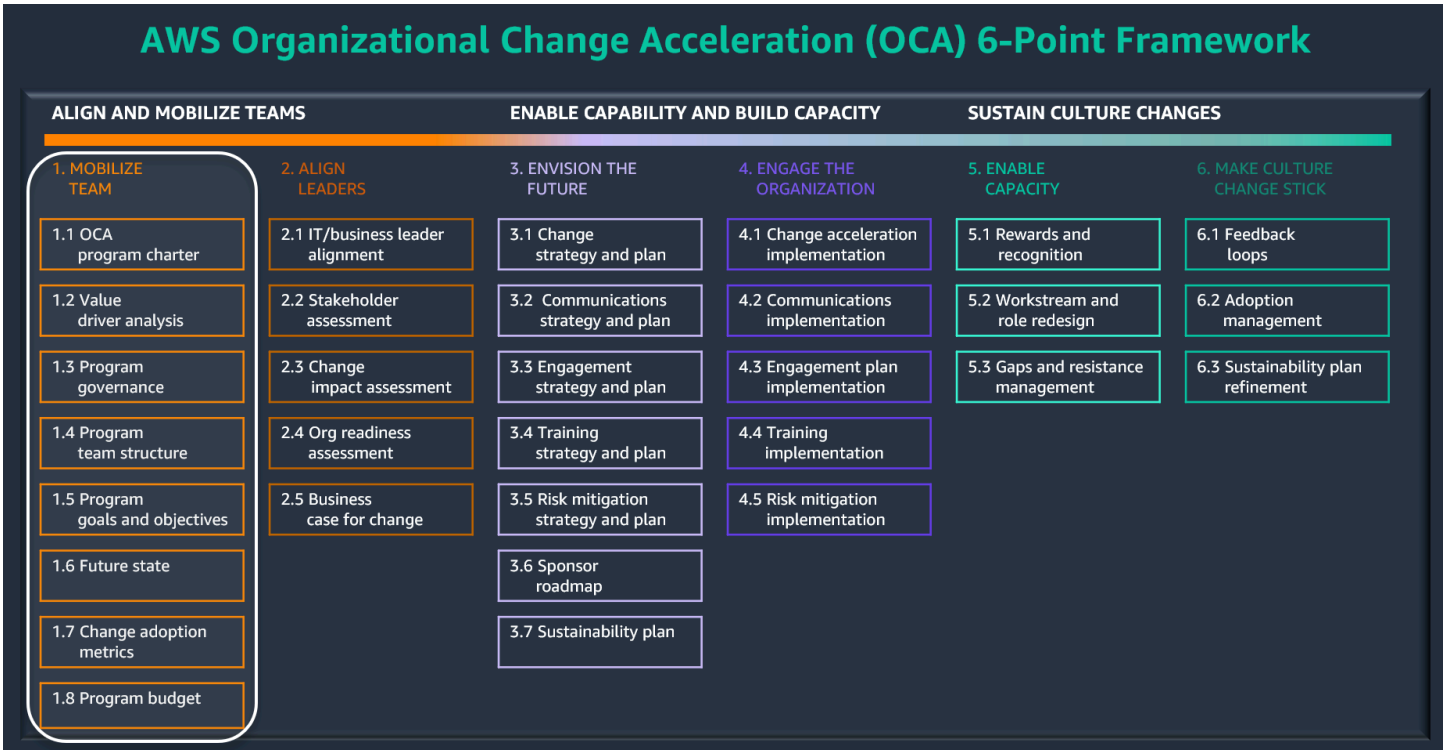

The first point, *Mobilize Team*, helps you start building structure and measures of success and governance around your change acceleration effort and activities. It contains eight subpoints:

- 1.1 OCA program charter. Describes the goals, sponsors, scope, structure, and processes of the change acceleration workstream, so you can lead the organization through the overall change acceleration effort effectively.
- 1.2 Value driver analysis. Helps you examine and define the specific paths to value creation.
- 1.3 Program governance. Ensures integrated alignment with executives, key stakeholders, the cloud program team, and the change acceleration team. It also covers ownership, decision rights, issue management, and escalation processes for change acceleration activities.
- 1.4 Program team structure. Ensures that the the organizational change acceleration team is staffed with dedicated and experienced organizational change resources whose roles are clearly defined in a responsible, accountable, consulted, informed (RACI) matrix.
- 1.5 Project goals and objectives. Establishes quantifiable cloud goals (both short-term and longer-term) that are aligned with strategic priorities.
- 1.6 Future state. Defines the organization's desired future state relative to culture, structure, people resources, technology, and processes.

- 1.7 Change adoption metrics. Provides a framework for measuring and documenting the progress and success of organizational change acceleration activities, and correlates change adoption metrics with desired business outcomes.
- 1.8 Program budget. Describes the costs that are associated with supporting the people, organizational, and upskilling dimensions of the organizational change acceleration program.

This guide discusses each subpoint of *Mobilize Team* 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 *Mobilize Team* phase of the AWS OCA 6-Point Framework contributes to the following outcomes:

- **Value realization and return on investment (ROI)**: Much of the value of cloud transformation depends on people. Implementing the subpoints in *Mobilize Team* helps prioritize and align people levers to the cloud strategy and desired business outcomes.
- **Transformational leadership**: Leadership is aligned and mobilized for accelerating cloud transformation.
- **Cloud acceleration**: The OCA workstream establishes the direction, metrics, governance, and program budget required to quickly mobilize resources to accelerate cloud transformation.
- **Organizational alignment**: The OCA workstream 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 (this guide)
2. Align Leaders
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.

1.1 OCA program charter

Overview

A formal *OCA Program Charter* document is intended to build leadership alignment and buy-in for the change acceleration scope of work from the onset of the cloud program. It specifies the cloud program team's dependencies on other areas and identifies key stakeholders. The OCA program charter includes:

- A review process for organizational change acceleration deliverables
- A definition of stakeholder responsibilities related to change acceleration activities
- A definition of change acceleration metrics and reporting requirements

The OCA program charter is purposeful, thoughtful, and structured. It delivers timely solutions and tactics to maximize speed, optimize adoption, and mitigate organizational risk. A cloud program inevitably includes people-related or political risks that might cause issues or derailments. Instead of reactively waiting for these issues, the charter document proactively anticipates and addresses risks by assigning deliverables, stakeholder roles and responsibilities, metrics, and reporting.

Best practices

Use the following tips to increase the effectiveness of your OCA program charter:

Scoping

The OCA program charter should be completed early in the cloud transformation project, as soon as a high-level cloud strategy and plan has been developed. This provides an early indication of the scope and breadth of OCA activities needed to accelerate cloud transformation.

- Clearly define what is in scope and out of scope for change acceleration activities. (For example, are the activities enterprise-wide or only for certain functional areas or verticals?)
- Confirm that program or project leadership understand the initial end-to-end OCA scope of work, dependencies, and key deliverables.
- Outline a formal review process for key deliverables.

- Establish a RACI that identifies key stakeholders that will be responsible, accountable, consulted, and informed of major change management activities or deliverables.
- Define initial OCA metrics and reporting requirements.

Assumptions

- The OCA program charter is developed and delivered by experienced organizational change resources. These resources could be provided internally, by a partner, or by AWS Professional Services.
- The OCA program charter is delivered as a component of a broader OCA solution.

Define responsibilities, requirements, resources, and results

The following table identifies the key people and data that the OCA program charter should involve and document.

Area	Responsible party	Task
Responsibilities	OCA lead	Draft the initial charter and iterate as necessary.
Requirements	Executive sponsor Customer leadership team Customer internal change team	Validate the charter, the data, and the outcomes it outlines.
Resources	Dedicated OCA lead Time with leadership team for validation and iteration	Identify the key resources needed to make the overall change acceleration engagement successful.
Results and KPIs	Executive sponsor and customer team (accountable for identifying the outcomes they want to see	Identify the results and key performance indicators (KPIs) that are agreed to, to measure organizational

Area	Responsible party	Task
	and the metrics they want to measure) OCA lead and customer internal change team (accountable for measuring the results and KPIs, and iterating on the plan)	readiness and the effectiveness of the change acceleration strategy and plan.

Define process

Step	Focus	Activities and people accountable for process
1	Supplier	AWS teams and customer teams
2	Inputs	• Business case • Scope • Timeline • Cloud plan or roadmap • Account plan • Program-level charter timelines or workplans by workstream • Desired business outcomes

Step	Focus	Activities and people accountable for process
		• Guiding principles • Strategic initiatives
3	Process	• Assess key program leaders to confirm expectations, identify areas of concern related to the change acceleration workstream, and understand change landscape. • Gather information related to program business case, timelines, schedules, and stakeholder information (if available). • Develop initial charter draft. • Review and validate initial charter draft. • Gather sign-off on final charter. • Implement the scope of work.

Step	Focus	Activities and people accountable for process
4	Outputs	• OCA program charter • Sign-off on charter
5	Customer	• Customer executive sponsors • Leadership teams • Customer change acceleration team • AWS leads and teams • All teams identified in RACI matrix

Develop charter document

Section	Heading	Subheadings (if applicable)
1	Overview	1.1 Change acceleration approach 1.2 Guiding principles in the transformation program 1.3. Key challenges 1.4 Key assumptions
2	Scope	2.1 Activity level scope

Section	Heading	Subheadings (if applicable)
3	Strategy and approach	3.1 Change acceleration journey 3.2 Key change acceleration activity and deliverable roles and responsibilities
4	Communications	4.2 Detailed communications planning 4.3 Case for change 4.4 Cultural observations and recommendations
5	Governance	5.1 Document repository and collaboration 5.2 Risk, action, issue, decision (RAID) management 5.3 Meeting cadence and status reporting 5.4 Change acceleration team structure and key roles and responsibilities 5.5 Change acceleration metrics 5.6 Stakeholder alignment and engagement 5.7 Change acceleration deliverable review and approval process

FAQ

Q. Who should be involved in the creation of the OCA program charter?

A. The required individuals were listed earlier in this section. In addition, the following people should be involved in developing the change acceleration charter: cloud leadership team, project or program leader, customer internal change team, internal service providers (for example, communications, training/learning, human resources) if they have a role in supporting the cloud initiative.

Q. When should the OCA program charter be used?

A. As soon as the OCA charter is developed, use it to plan OCA resourcing, timelines, and high-level OCA activities in support of the cloud strategy. Review the charter with all impacted stakeholder groups.

Q. What is the output of the charter?

A. The output of this exercise is an approved and aligned OCA program charter with assigned deliverables, outlined stakeholder roles and responsibilities, metrics, and reporting. The OCA program charter also informs 3.1 Change strategy and plan in this framework.

Additional steps

To begin building an OCA program charter, follow these steps:

1. Assess key program leaders to confirm expectations, identify the areas of concern related to the change acceleration workstream, and understand the change landscape.
2. Gather information about the program business case, timelines, schedules, and stakeholder information (if available).
   a. Meet with program leaders to gather information about the business case, scope, timeline, milestones, level of effort, and impacted stakeholders.
   b. Meet with the executive sponsor to gather information about the cloud strategy and desired business outcomes, and to establish expectations for active and visible sponsorship.
   c. Meet with workstream leads to gather information about the scope, timing of critical deliverables and events, and expectations for interaction with the change acceleration team.
   d. Meet with internal groups such as organizational change management, corporate or strategic communications, employee engagement, human resources, training/learning, and

development as applicable to understand the level of support they will dedicate to the cloud program, and any expectations for change acceleration reports that you will need to provide to them.

3. Develop an initial charter draft.
4. Review and validate the initial charter draft.
5. Gather sign-off on the final charter.

1.2 Value driver analysis

Overview

Maximizing shareholder value is an important corporate objective, but it isn't specific and accountable enough for leaders who must also know which factors influence value the most and which factors can be most easily affected. These factors are known as *value drivers*, and they are the primary focus of organizations that succeed in maximizing shareholder value. Organizations can identify key drivers of value creation and structure a performance measurement approach around these value drivers.

Value driver analysis is an important foundation for strategic planning and helping management sort through their operations to define critical strategic levers. A value driver analysis presents an approach to increasing performance that will forge stronger links between operating performance measures and shareholder value creation. Value drivers can be categorized as growth drivers, efficiency drivers, or financial drivers. Companies tend to create paths to value creation by investing in growth opportunities, investing in operating efficiency, divesting from value-destroying activities, and reducing capital costs. Some organizations manage their business as if every operating factor were equally important.

Use value driver analysis when you want to examine and define the specific paths to value creation by function and level within the organization. This will help managers focus their attention on factors that matter the most. Value drivers should have a significant impact on maximizing shareholder value, and should be controllable.

The following value driver analysis matrix shows the correlation between management influence and value impact.

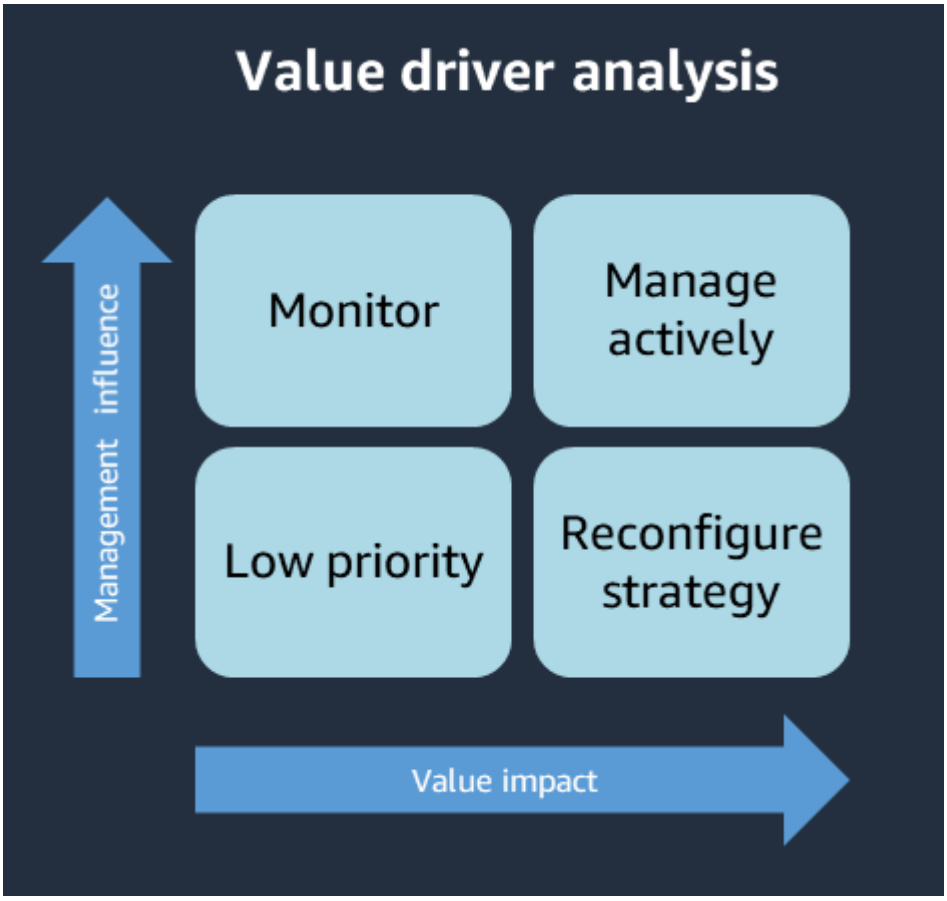

As the matrix shows:

- Value drivers that have a high impact to value and high degree of management influence should be managed actively.
- Value drivers that have a high impact to value and a low degree of management influence should be reconfigured by changing the strategy.
- Value drivers that have a low impact to value and a high degree of management influence should be monitored.
- Value drivers that have a low impact to value and a low degree of management influence should be considered low priority.

Best practices

Identifying and managing value drivers help the leadership team focus their attention on activities that will have the greatest impact on maximizing shareholder value. This focus enables

management to translate the broad goal of value creation into the specific actions most likely to deliver that value.

There are three categories of value drivers: growth drivers, efficiency drivers, and financial drivers. As shown in the following diagram, companies tend to manage these value drivers in four ways: investing in value-creating growth opportunities, investing in operating efficiency, divesting from value eroding activities, and reducing capital cost. By focusing on value drivers, management can prioritize the specific activities that will affect performance in each area.

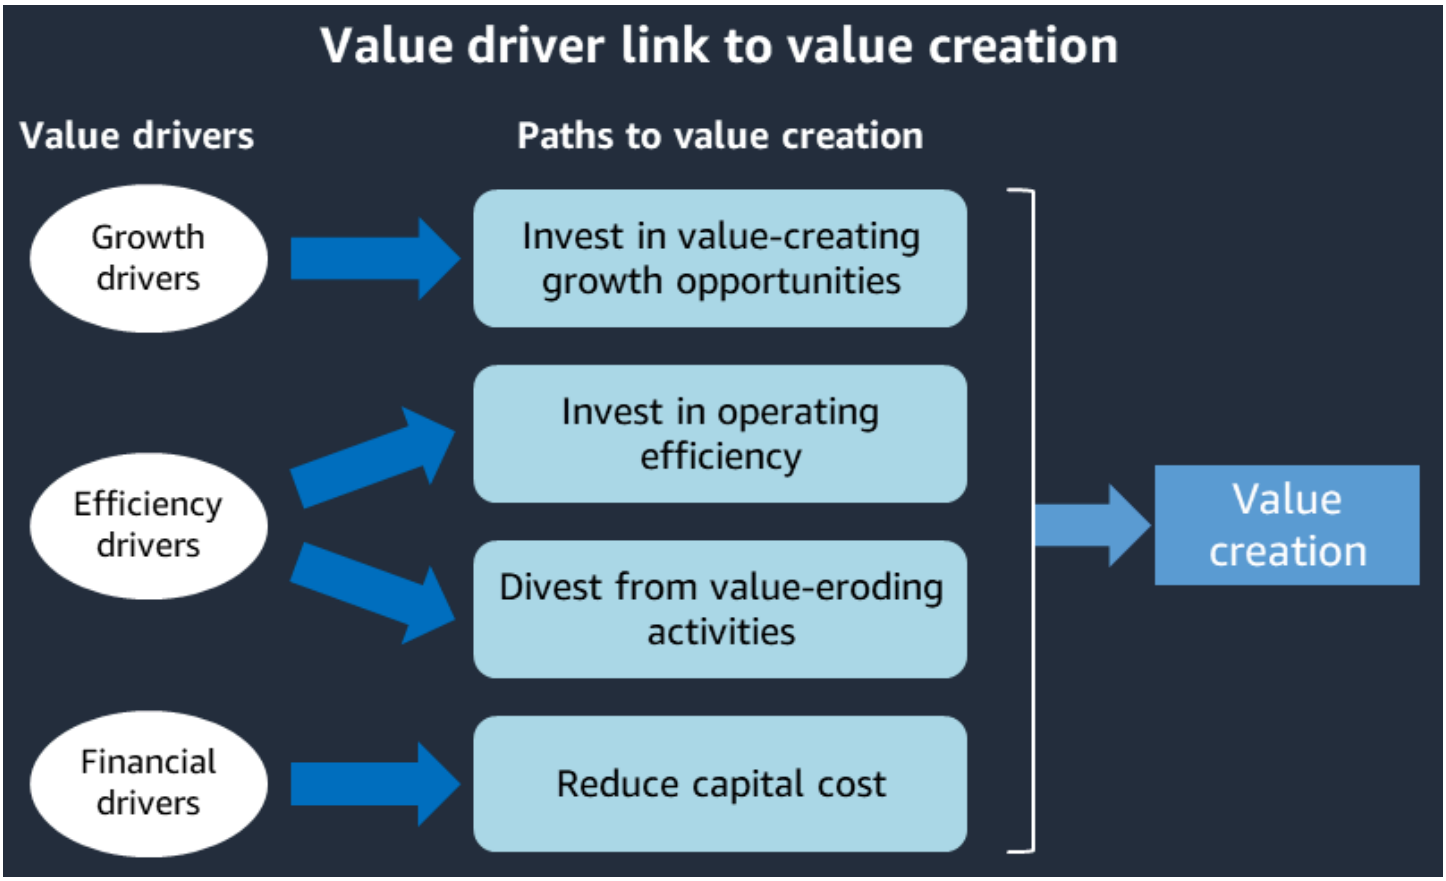

Examining and defining paths to value creation enables companies to identify and understand responsibilities by function and level within the organization. This, in turn, helps managers focus their attention on factors that really matter.

Often, IT organizations manage their business by treating every operating factor as equally important. IT managers have a solid knowledge of the variables that impact business performance and they manage that list aggressively. However, the list of variables is often too long and might be prioritized against goals other than value creation. Valuable resources are rationalized to increase market share, maintain pricing, increase distribution, introduce new products, increase operating efficiency, and so on without a clear sense of what true value drivers are.

Early in the cloud journey, organizations should explicitly define the value that they want to derive from the cloud. The cloud can positively impact all three value drivers (growth, efficiency, and financial value). A common best practice is to develop value maps that define all cloud initiatives that will impact value drivers, as shown in the following example.

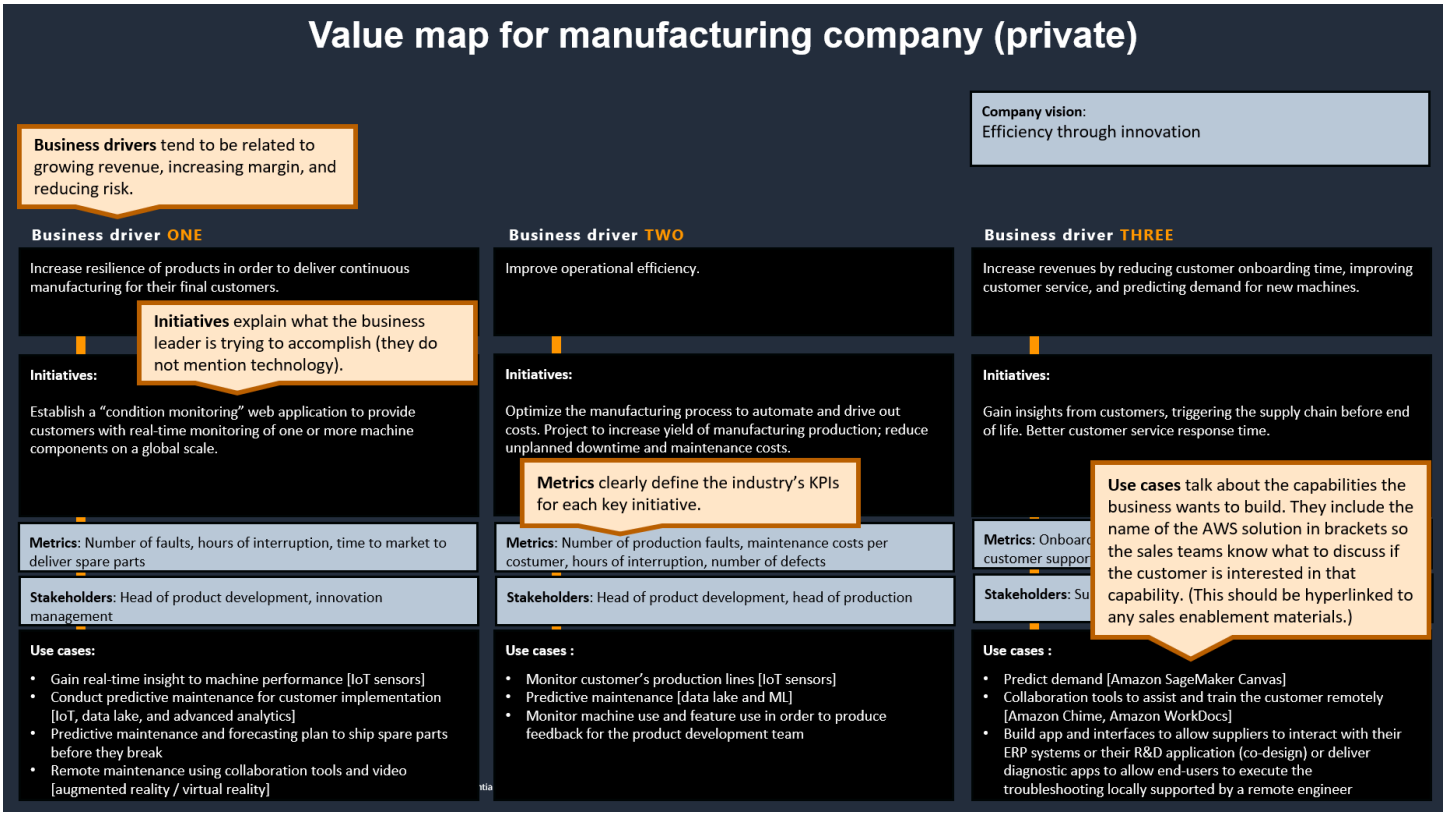

Value maps include the following information:

- Business/value driver: The penultimate description of business value. These tend to be financial measures relating to increased revenue, decreased cost, improved margins, and so on.
- Initiatives: What the business leader is trying to accomplish. Initiatives don't include any references to technology.
- Metrics: The measures used to quantify the success of the initiative over time.
- Use cases: The capabilities that the business wants to build to enable the initiative. The use case describes the technology that is used to establish capabilities.

Value maps should be created based on the organization's strategic priorities. If the key component of the strategy is revenue generation, make sure that your value maps address revenue-generating cloud initiatives, but do not exclude value maps for efficiency and financial drivers. This will allow

for a more complete picture of the value generated by the cloud and create momentum for continued cloud adoption.

FAQ

Q. Why is this analysis valuable?

A. Organizations can identify key drivers of value creation and structure a performance measurement approach around them. Leaders can, in turn, focus their attention on activities that have the greatest impact on value.

Q. When do you use it?

A. Use value driver analysis early in the cloud journey to determine how the cloud can influence growth, efficiency, and financial drivers. Use value driver analysis to develop more a detailed business case for the cloud.

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 analysis?

A. Value driver analysis uses discovery assessment outputs, external benchmarking, and strategic plans as inputs.

Q. What are the outputs of this analysis?

A. The analysis produces two outputs: the value driver matrix and value maps. The value driver matrix helps you understand your organization's value drivers by breaking down the broad operating parameters of the business into progressively smaller components until you reach the level where daily operating management decisions are made. The matrix also helps document which specific factors influence broad measures such as sales growth, operating profit, and so on. Value maps connect value drivers and business outcomes to specific cloud initiatives and use cases.

Additional steps

To develop a value driver map of your business or initiative, follow these steps:

1. Review strategic plans against cloud use cases that can impact value drivers. In other words, always ask the question: *Are we extracting the maximum value (revenue, efficiency, and financial value) from the cloud relative to our strategic priorities?*
2. Develop value maps for productivity, efficiency, and financial value drivers.
3. Develop and refine your cloud strategy to define the specific cloud initiatives, use cases, and metrics that contribute to each value driver.
4. Socialize value maps and cloud strategy with cross-functional leaders and middle management. Middle management typically plays a pivotal role in this activity, because they lead the largest number of employees and have to split their time between strategy and execution.
5. Develop a measurement plan to demonstrate the effects of executing against cloud use cases (leading indicators) on value drivers (lagging indicators).

1.3 Program governance

Overview

Program governance secures integrated alignment with executives, key stakeholders, the cloud program team, and the OCA team. It also defines ownership, decision rights, issue management, and the escalation process for change acceleration activities.

Use governance to achieve the following:

- Enable clear, decisive leadership and accountability of the change acceleration program.
- Provide quality assurance and a path to escalate issues and risks.
- Specify a decision rights framework for the program.
- Align workstream structure to the existing project, program, and organization-wide governance structure.
- Establish a cadence of meetings and scrum ceremonies that map to the rest of the program rhythm and reporting mechanisms.

Best practices

A program that establishes a clear structure and governance has a higher probability of success compared with one that does not. This is because decisions and ownership around decision-making rights can often be a major factor in delaying cloud programs. Governance establishes decision-making authorities, and can provide guidance around two-way door decisions (which can be made quickly with low risk and can be easily reversed), and one-way door decisions (which require more thought and contemplation, because they cannot be easy reversed).

For effective program governance, follow the best practices described in this section.

Establish organizational change acceleration team structure

Establish a high-level structure that depicts reporting relationships and high-level responsibilities, as shown in the following example.

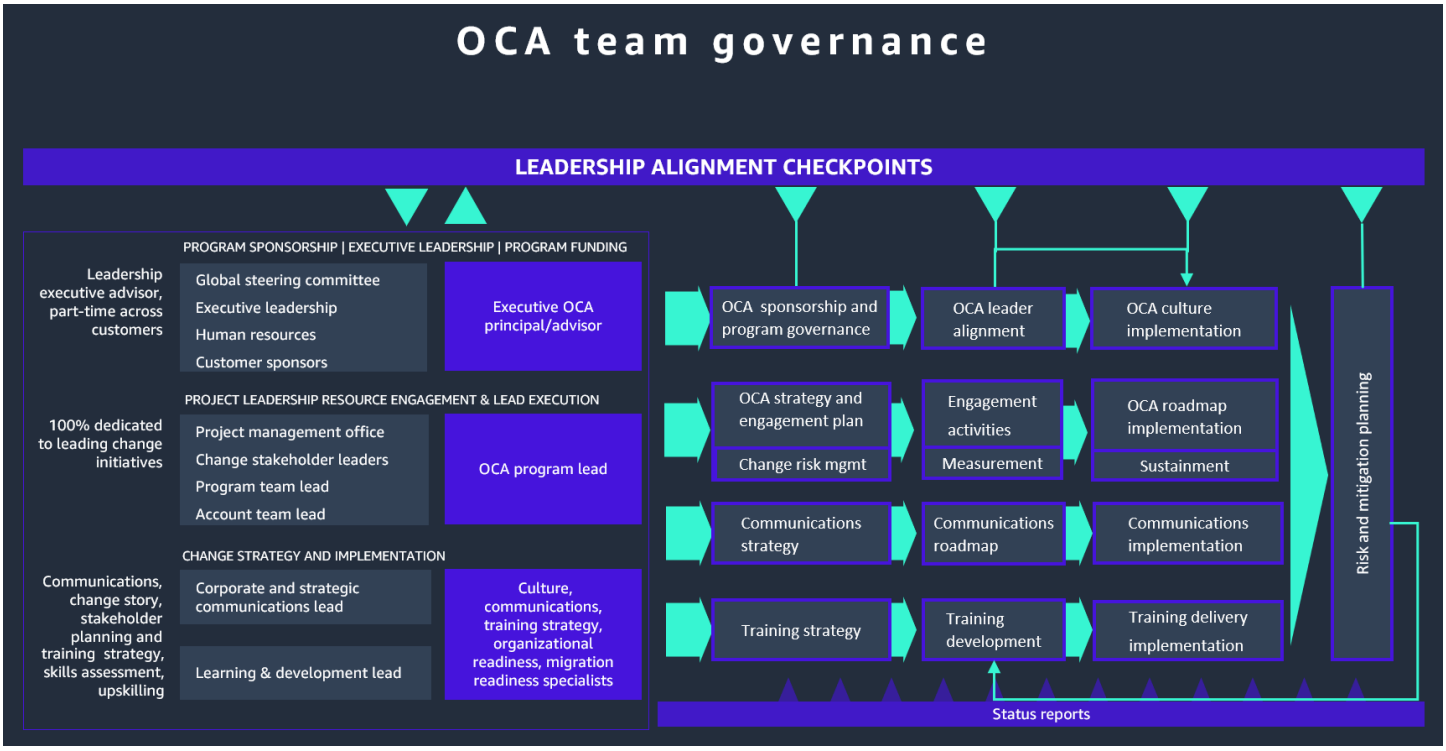

Define roles for the OCA team

Define the roles of the organizational change acceleration team. Provide the high-level responsibilities for each role, as shown in the following example.

OCA team role descriptions

OCA necessary roles	
This workstream takes the *two-in-the-box* approach for delivery. The goal of this approach is to enable knowledge transfer between the change acceleration SMEs and customer leaders for rapid deployment and sustainable change ownership. In the end, this leaves customers better off.	
Executive OCA principal/advisor	**OCA program/workstream lead**
This role works with the PMO, and secures participation from sponsors and executive leaders to align executive stakeholders, to coach leaders on culture change leadership responsibilities, and to create cloud leadership capabilities, partnering with the executive sponsor to enable change impact decision-making.	This role leads the day-to-day program with deep domain expertise. It manages workstream strategy and execution of deliverables, develops metrics and tracking to ensure program success, approves and validates change strategy and roadmap, and manages and tracks risks.

Additional OCA SMEs		
Depending on the scale and complexity of the AWS module, more roles are usually needed, and these require additional layers of expertise in communications, training, employee module, organizational change readiness, infrastructure/migration readiness, and CCoE. They could be full-time or part-time, based on scale and complexity .		
Communications specialist	**Training specialist**	**Culture, workforce, leadership specialists**
This role is focused on building out detailed communications for all stakeholder groups. This role is needed if scale and complexity of the engagement exceeds capacity of work for base roles, and/or if a deeper area of depth is needed. Ideally, a customer stakeholder is identified for this role.	This role is focused on building out the detailed training strategy and roadmap. This role is needed if the scale and complexity of the engagement exceeds the capacity of work for base roles, and/or if a deeper area of depth is needed.	These roles are focused on culture, workforce, or leadership backlog activities. These roles are needed only if the scale and complexity of the engagement requires a deeper area of depth in any of these domain areas.

These roles and responsibilities might include the following.

OCA necessary roles:

- Executive OCA principal/advisor: Works with the PMO and secures participation from sponsors and executive leaders to align executive stakeholders, to coach leaders on culture change leadership responsibilities, and to create cloud leadership capabilities, partnering with the executive sponsor to enable change impact decision-making.
- OCA program/workstream lead: Leads the day-to-day program with deep domain expertise. Manages workstream strategy and execution of deliverables, develops metrics and tracking to ensure program success, approves and validates change strategy and roadmap, and manages and tracks risks.

Additional OCA SMEs:

- Communications specialist: Builds out detailed communications for all stakeholder groups. This role is needed if scale and complexity of the engagement exceeds capacity of work for base roles, or if a deeper area of depth is needed. Ideally, a customer stakeholder is identified for this role.

- Training specialist: Develops the detailed training strategy and roadmap. This role is needed if the scale and complexity of the engagement exceeds the capacity of work for base roles, or if a deeper area of depth is needed.
- Culture, workforce, leadership specialists: These roles focus on culture, workforce, or leadership backlog activities. They are needed only if the scale and complexity of the engagement require a deeper area of depth in any of these domain areas.

Define the OCA roles in your organization

Define the high-level responsibilities of relevant OCA roles in your organization, as in the following example.

OCA customer role descriptions

CUSTOMER RESOURCE AVAILABLITY		
EXECUTIVE SPONSOR 25% PROJECT ALLOCATION	**ORG CHANGE ACCELERATION LEAD** 100% PROJECT ALLOCATION	**PROGRAM/ENGAGEMENT LEAD(S)** 100% PROJECT ALLOCATION
• Member of steering committee • Executive role who makes decisions on significant issues • Partner with OCA advisor on change strategy and decision-making • Review and approve scope changes • Review and approve project and program plans • Provide final project approval	• Partner with OCA lead on co-development and delivery of change assessments, strategy, and roadmap • Assist to identify stakeholder groups, business objectives, risks, blockers, and synergies • Provide active and visible change leadership • Serve as organization change champion	• Assist in workshop recording and documentation • Assist in scheduling assessment interviews • Serve as first point of escalation • Review and comment on deliverables • Review and comment on project plans
LINE OF BUSINESS OWNER(S) 30% PROJECT ALLOCATION	**TRAINING AND COMMUNICATIONS SMEs** 30% PROJECT ALLOCATION	**HR BUSINESS PARTNER** 20% PROJECT ALLOCATION
• Work with OCA team to define business needs • Attend workshops and interviews • Provide subject matter expertise for portions of the project • Review and comment on deliverables	• Co-develop with OCA messaging to create awareness and engagement • Provide messaging standards and guidelines • Provide translation support • Partner with OCA team on communications strategy and roadmap	• Partner with OCA team to provide targeted HR support at point of need (before and after transition)

These roles and time commitments might include the following.

Executive sponsor (25%)

- Member of steering committee
- Executive role who makes decisions on significant issues
- Partner with OCA advisor on change strategy and decision-making

- Review and approve scope changes
- Review and approve project and program plans
- Provide final project approval

Organizational change acceleration lead (100%)

- Partner with OCA lead on co-development and delivery of change assessments, strategy, and roadmap
- Assist to identify stakeholder groups, business objectives, risks, blockers, and synergies
- Provide active and visible change leadership
- Serve as organization change champion

Program/engagement lead(s) (100%)

- Assist in workshop recording and documentation
- Assist in scheduling assessment interviews
- Serve as first point of escalation
- Review and comment on deliverables
- Review and comment on project plans

Line of business owner(s) (50%)

- Work with OCA team to define business needs
- Attend workshops and interviews
- Provide subject matter expertise for portions of the project
- Review and comment on deliverables

Training and communications SMEs (30%)

- Co-develop with OCA messaging to create awareness and engagement
- Provide messaging standards and guidelines
- Provide translation support
- Partner with OCA team on communications strategy and roadmap

HR business partner (20%)

- Partner with OCA team to provide targeted HR support at point of need (before and after transition)

Define risk identification and mitigation strategy

Develop a strategy for the ongoing identification of people and change risks that have the potential to impede the cloud strategy and value realization. This strategy should have the capability for broad-based inputs of potential risks as well as agile mechanisms for risk analysis, planning, mitigation, and tracking. The following diagram provides an example.

The objective of the risk mitigation strategy and plan is to ensure a seamless process to manage status, issues, and escalations and resolve conflicts in a timely manner.

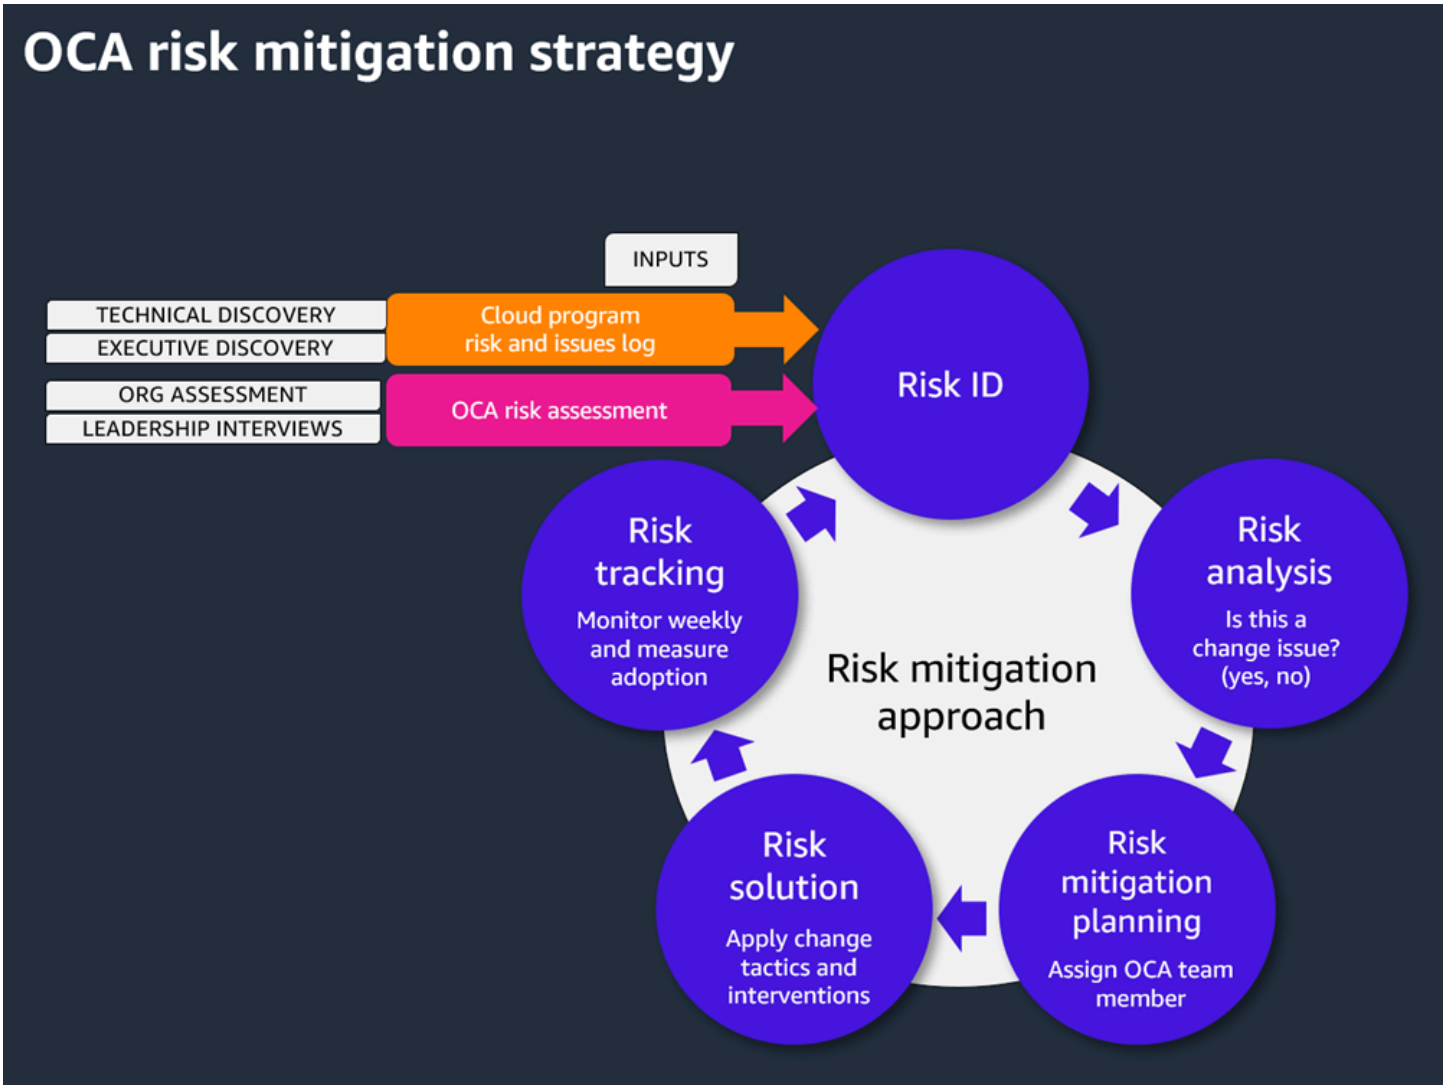

The change risk assessments guide the change activities and detailed change plan to help ensure that change solutions are targeted where they are needed and where they will drive the most impact. Targeted change solutions enable OCA to influence the right audiences, at the right times, in the right manner.

Develop a RACI for OCA

Lastly, define the specific accountabilities and responsibilities for key organizational change acceleration capabilities for both the customer and AWS. Here's an example responsible, accountable, consulted, informed (RACI) matrix.

R - Responsible A - Accountable C - Consulted I - Informed	Organizational Change Acceleration (OCA) Cloud Program RACI							
	Partner OCA Lead (AWS or other consultants)	Partner Training Lead (AWS or other consultants)	Customer OCA Lead	Partner Cloud Program Workstream Lead(s) (AWS or other consultants)	Customer Cloud Program Workstream Lead(s)	Customer Cloud Transformation Program Leader	Customer Cloud Executive Sponsor	Additional Representatives (e.g., HR, BU Leads, etc.)
Organizational Change Acceleration								
Organizational Readiness	R	I	A	C	C	C	I	C
Sponsorship, Leadership, & Culture	R	I	A	C	C	C	I	C
Strategy, Roadmap, Execution	R	R	A	C	C	I	I	C/I
Communications	C/I	C/I	A/R	I	I	C/I	I	I
Organizational Alignment	R	I	R	I	I	A	I	I
KPIs	C	C	R	I	I	A	I	I
Program Communications (Email)								
Draft Communications (Structure)	R	I	A	C	C	C	I	C
Draft Communications (Detail)	R	I	A	C	C	C		C
Content Review	R	R	A	C	C	I		C/I
Content Approval	I	I	I	I	I	A		I
Final Communication Approval	I	I	I	I	I	A	I	I
Final Communication Sent	I	I	A	I	I	I		I
Migration and/or Wave Communications (Email)								
Draft Communications (Structure)	R	I	R	C	C	C		C
Draft Communications (Detail)	R	I	R	C	C	C		C
Content Review	R	R	A	C	C	C		C
Content Approval	I	I	I	I	I	C		I
Final Communication Approval	I	I	I	I	I	C	I	I
Final Communication Sent	I	I	I	I	I	C		I
Communication Site (SharePoint, etc.)								
Design Wireframe for Site	R	R	A/R	C	C	C		C
Approve Site Contents	C	C	R	C	C	A	I	C
Build out site	C/I	C/I	A/R	I	I	I		I
Create and/or Upload Site Content	R	R	A	I	C/I	C/I		C/I
Make iterative updates to site	C	C	A/R	I	I	C		C/I

FAQ

Q. Who should be involved in program governance?

A. Executive sponsors, project or program leader, change leader, internal service providers (for example, communications, training/learning, and human resources, if they have a role in supporting the change initiative).

Q. What are the inputs to this activity?

A. A project charter, business case, input from the steering committee, cloud leadership team, and any cloud assessment results (for example, migration readiness or cloud maturity assessments).

Q. What are the outputs of this activity?

A. A status reporting template, reporting processes, project team organization charts, decision matrix, and RACI.

Q. When should the activity be conducted?

A. The OCA governance structure must be established at the beginning of the initiative and remain in place until project completion. As the team expands or contracts, new roles should be defined and approved by program leadership. When team changes occur, the organizational charts should be updated and those changes should be communicated across the program.

Additional steps

1. Review the cloud governance structure and processes such as Cloud Center of Excellence (CCoE), program management office, or transformation management office.
2. Determine the integration of OCA into cloud governance structures and processes.
3. Develop the OCA strategy, process, and tools for risk management.
4. Integrate the OCA risk management process with the program-level risk management process.
5. Develop status reporting processes.
6. Develop standardized OCA status templates for reporting purposes.

1.4 Program team structure

Overview

Maximizing the value of cloud requires change acceleration support to address the *people* aspects of the change. By staffing your team with dedicated and experienced organizational change management resources, you can get much closer to enabling a smoother transition to the cloud.

The change acceleration team operates within the cloud transformation team and is responsible for identifying organizational changes, defining role changes, planning communications, mapping training requirements, and securing executive sponsorship.

Best practices

Staffing a strong and effect OCA team at the beginning of the program is critical. Evaluate staffing levels on an ongoing basis to determine whether they should be scaled up or down in connection with the program's scope and timeline.

Here are some example key roles and responsibilities in the project team:

- OCA executive advisor: Engages with the executive program sponsor and other IT and business leaders who are responsible for cloud transformation (for example, CIO, CTO, cloud program director, CCoE leader).
- OCA lead: Manages all aspects of the change acceleration team, deliverables, and timelines at the program level. Works with customer workstream counterpart, program manager, cloud program director, CCoE leader, and other program workstream leads.
- Executive change acceleration oversight and program oversight roles: Collaborate at all levels to drive project strategy and successful implementation with responsibility for quality assurance.
- Communications lead: Establishes the communication strategy and implements the communication plan; works with the customer communications lead and other stakeholders, such as business leads and application owners, as required.
- Training lead: Designs and develops the training strategy and plan. Works in collaboration with the learning and development or training lead to determine how to best advertise training, target users for training courses, handle training logistics, and roll out training within the customer's environment.

- Specialty subject matter experts (as needed): Focus on variable aspects of the program such as culture analysis, diversity and inclusion, and strategic workforce planning.

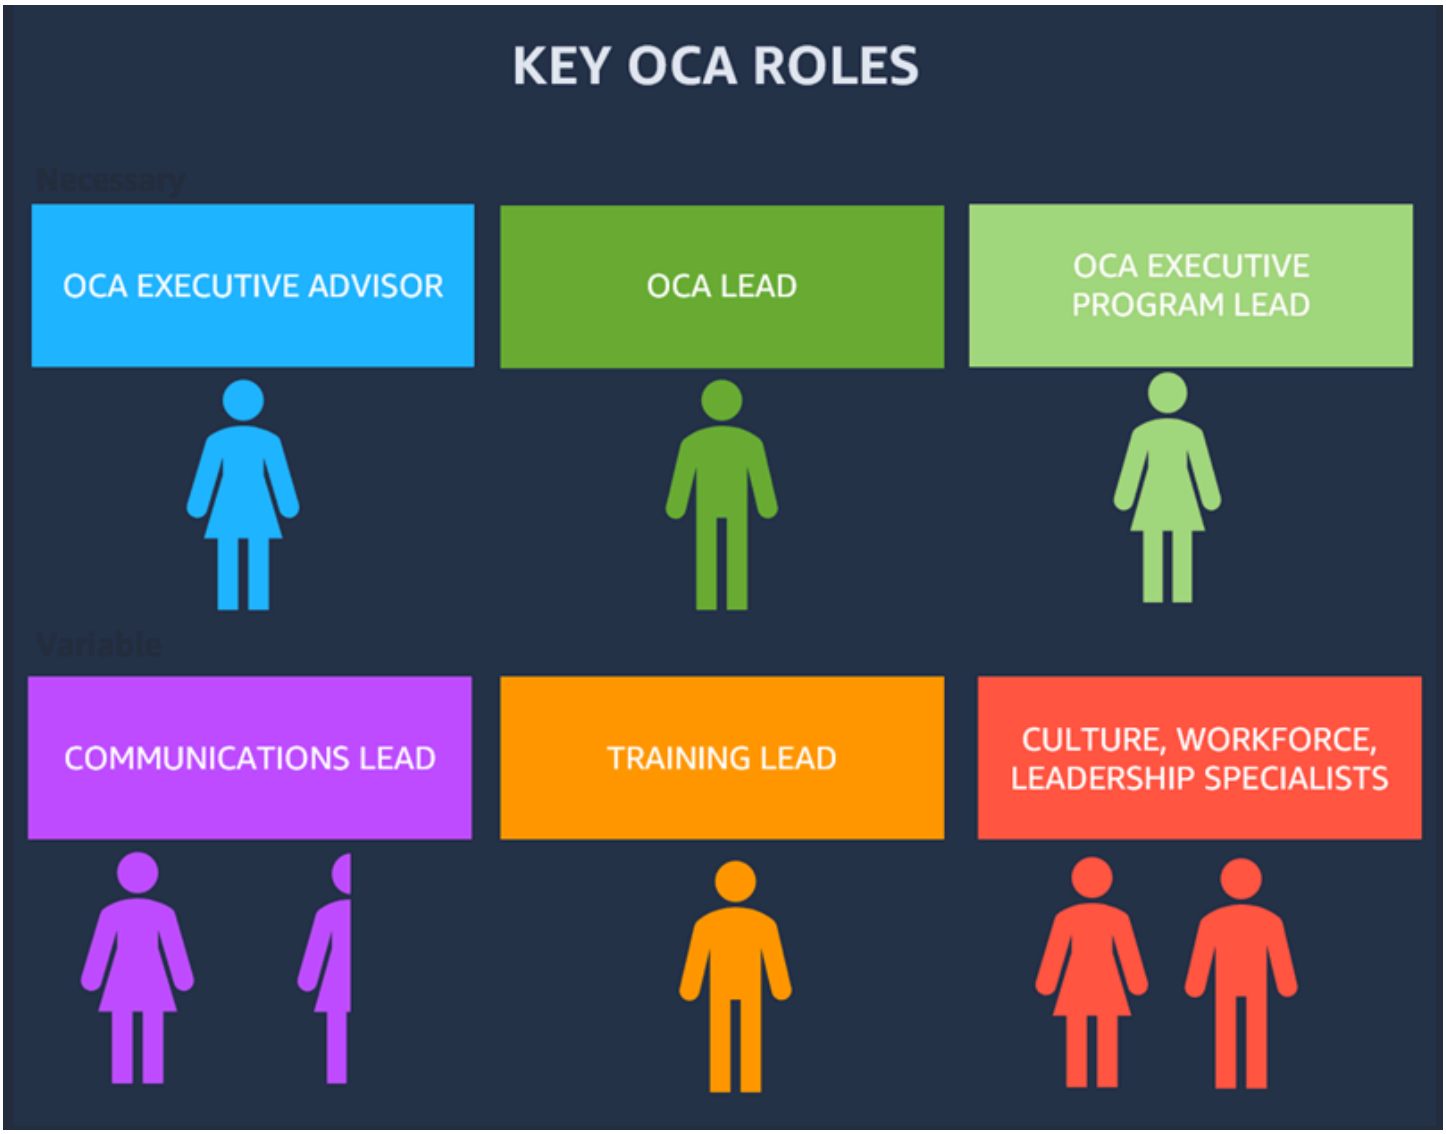

For a majority of projects, three roles are key: people transformation or change acceleration lead, organizational readiness and communications lead, and training lead. These three roles are the foundation for the change acceleration team that supports a cloud transformation program. Additional resources can be added to the team as the scope of cloud transformation increases or timelines shift.

FAQ

Q. Who should be involved in this activity?

A. Executive sponsors, project or program leader, change leader, internal service providers (for example, communications, training/learning, and human resources, if they have a role in supporting the change initiative).

Q. What are the inputs to this activity?

A. Inputs might include the OCA program charter, outputs from cloud readiness assessments such as the AWS Migration Readiness Assessment (MRA), cloud strategy and plan, and discovery materials.

Q. What are the outputs from this activity?

A. Project team organizational structure and OCA team roles and responsibilities.

Q. When should this activity be conducted?

A. The OCA team structure should be built and staffed as soon as program governance is defined.

Additional steps

To create the project team structure, complete the following:

1. Review the OCA program charter.
2. Review the scope and cloud readiness assessment outputs.
3. Review the change impact assessment outputs.
4. Review the high-level change Impacts (if available).
5. Review the discovery materials (as needed).
6. Engage cloud project leadership, internal change team (if available), HR business partner, and internal communications or training leaders for input on team structure.
7. Interview IT functional leaders to validate gaps, organizational changes, and impacts to roles.
8. Confirm and validate the work effort required for each role:
   - OCA lead
   - Communications lead

- Training lead
- Other roles as necessary

9. Present the team structure to cloud leadership for approval.
10. Identify internal resources or retain staffing externally.
11. Onboard the OCA team.
12. Introduce, onboard, and engage the OCA team.

1.5 Program goals and objectives

Overview

The definition of cloud goals and objectives originates in the discovery phase and are often refined during cloud readiness assessments (such as the MRA) and cloud planning (such as cloud use case prioritization, migration planning, and cloud roadmap) efforts. Use the business case, interviews, and strategy documents to articulate a clear, concise, and compelling set of goals and objectives. These goals and objectives should be specific, measurable, aligned, and time-bound. Furthermore, the goals should adequately represent business and IT. When cloud goals and objectives are documented and agreed upon, they can be used to increase awareness, understanding, and adoption.

Best practices

- Include multiple stakeholder groups and perspectives when defining project goals and objectives. Include:
  - Business leadership
  - IT leadership
  - External customers
  - Employees
- Refine the goals from the business case and OCA program charter to ensure that they are tangible, concrete, measurable, and manageable targets that represent planned progress toward the adoption of the future state.
- Use cloud goals to drive greater alignment between business and IT.
- Use cloud goals to establish prioritization of cloud activities such as migration, modernization, culture, and ways of working.
- Use cloud goals as a motivational tool to encourage people to achieve high levels of performance, and as a basis for celebration, reinforcement, and rewards. For additional guidance, see 5.1 Rewards and recognition in this framework.

FAQ

Q. What is it?

A. Cloud goals and objectives originate in the discovery phase and are refined during the assess and planning phases through mechanisms such as the Migration Readiness Assessment (MRA), Migration Readiness Planning (MRP), cloud use case definition and prioritization, business value maps, and cloud strategy/roadmap. The OCA team aligns its activities against those goals and objectives, and embeds them in the strategy. Goals and objectives are based on the business case, customer interviews, strategic plans, and MRA and MRP findings.

Q. Why is it valuable?

A. Including the change OCA team in assessment and planning sessions builds alignment among the people, process, and technology aspects of migrating and modernizing applications and workloads on AWS. When cloud goals and objectives are documented and agreed upon, they can be used to increase awareness, understanding, and adoption. Cascading cloud-related goals provide direction, clarity, and focus for daily behaviors. The cascading cloud goals send signals about the relative priority of the cloud and creates aligned actions that contribute to successful cloud transformation.

Q. When do you use it?

A. Use project goals and objectives to motivate, monitor, and measure progress on the cloud adoption journey. First understand which goals have already been established. Then work to establish new goals that are focused and simple. If the goal isn't easily understood, it's probably not the right goal. Build metrics and measurement mechanisms to update business leaders on the progress against these goals, and forecast business scenarios based on new implications. The project goals and objectives can be developed and implemented in conjunction with the case for change and leadership action plans as a part of the OCA 6-Point Framework.

Q. Who should be involved in this activity?

A. Executive sponsors, project or program leader, change leader, internal service providers (for example, communications, training/learning, and human resources, if they have a role in supporting the change initiative).

Q. What are the inputs to this activity?

A. Business case, discovery phase outputs (MRA and MRP), interviews with the executive sponsor and human resources, cloud strategy, and business value realization plans.

Q. What are the outputs of this activity?

A. Documented goals and objectives for the project, a measurement and monitoring plan, and an initial plan for cascading goal communications.

Q. What is the right number of goals and objectives? How many is too many?

A. As the number of activities or results that are measured increases, employees become distracted and their efforts are diluted. Therefore, we recommend that you focus on the few critical areas of performance.

Q. What are common cloud metrics that could be used for goals and objectives?

A. Metrics include:

Cost savings

- IT spend on application per user
- Total IT infrastructure spend

Staff productivity

- Virtual machines (VMs) managed per administrator
- TBs managed per administrator

Operational resilience

- Application availability
- Total monthly incidents
- Critical (P1/P0) incidents
- Security incidents
- Mean time to recovery (MTTR)
- Application resilience rating
- Application security rating

Business agility

- New products deployed
- Time to market
- Time to deployment
- Code deployment frequency
- Customer satisfaction

Q. How can project goals and objectives be used as a part of an OKR strategy?

A. Objective and key results (OKRs) consist of an *objective* (a significant, concrete, clearly defined goal) and three to five *key results* (measurable success criteria used to track the achievement of that goal). Depending on the objective, cloud metrics can be transformed into key results statements such as the following:

- **Objective**: Improve the customer experience.
- **Key result**: Increase the number of new products deployed 100 percent over the next 12 months.

Additional steps

1. Decompose goals and objectives into more granular and specific objectives. Objectives that are measurable allow teams to track progress, understand if needs were addressed, and know whether a change was effective. Use SMART criteria to guide the definition of objectives:
   - Specific: The objective has an observable outcome.
   - Measurable: You can quantify or indicate progress on the outcome.
   - Achievable: The outcome is realistic and feasible.
   - Relevant: The objective aligns with or supports other goals or strategic initiatives.
   - Time-bound: You can set a target date for the effort.
2. Describe mandatory design elements, prescribed implementation details, or aspects of the current state and planned future state that should not be changed by the solution. Constraints are limitations that can be addressed when proposing alternative options. Some examples of constraints are:
   - Budgetary restrictions
   - Time restrictions

- Technology
- Infrastructure
- Policies
- Limits on available resources
- Restrictions based on the skills of the team and stakeholders
- A requirement that certain stakeholders shouldn't be affected by the solution
- Compliance with regulations

3. Describe the beliefs that determine whether the future state meets business needs. In an uncertain environment, it can be difficult to prove that a planned change will meet a business need. Assumptions are defined so that appropriate course corrections can be made, including redirection or termination of the initiative if an assumption proves invalid.
4. Review the alternatives considered in the business case and determine whether there is flexibility to evaluate other options. If so, indicate what kind of options will and won't be considered when investigating possible solutions, including changes to the organizational structure or culture, capabilities and processes, technology and infrastructure, policies, products, or services.
5. Identify the potential value of the solution, which is the net benefit of the solution after accounting for operating costs. In general, a change must result in greater value for the organization compared with no action being taken. In some cases, the future state presents a decrease in value from the current state. For example, responding to increased competition or complying with new regulations decreases the overall value but is necessary to remain operational. Express the potential value in terms of the expected benefits, expected costs, and likely result if no change is made.
6. Refresh the change acceleration strategy and plan (OCA 3.1) as necessary.
7. Communicate cloud goals and objectives to leaders across the organization, and revise them based on input.
8. Cascade cloud goals and objectives down organizational levels to the front line.
9. Link cloud goals to individual performance review processes or HR systems.
10. Communicate successes to all employees and link the project to overall company goals.

1.6 Future state

Overview

A cloud future state identifies the vision and potential value that can be realized from the cloud solution. The cloud future state is derived from organizational assessments, external benchmarking, and cloud strategy. It represents a visual alignment of the organization culture, structure, people, technology, and process design to the new, cloud-centric ways of working.

The desired cloud future state informs your approach to transforming your people, skills, and organization. Some analysis techniques that can guide the definition of the future state are decision analysis, process analysis, business capability analysis, feature decomposition, prototyping, and product roadmapping. The future state is heavily dependent on, and should be aligned to, the cloud strategy.

Best practices

The future state should align with cloud strategy overall: What benefits will the cloud bring to the organization and to the people within it? What is the value that will be generated that was discussed in the strategy? These key insights from the strategy are part of the foundational building blocks that define the future state. Many companies benefit from identifying a network of change agents that represent a footprint of the impacted user base (functions, geographies, roles, and so on). A change agent is someone who is knowledgeable, authentic, and credible, and has influence within their network, even if they don't have formal authority.

Think about organizational alignment and establish an ongoing partnership between organizational structures, business operations, talent, and culture. The future will look different for every organization, but you can follow these three steps to help define your future state.

Step 1. Gather necessary information

Culture	Structure	Processes
How will people behave in the future state?	How should the organization be organized?	What processes are key to the organization's vision?

Culture	Structure	Processes
What will they believe is important?	How much management should there be?	How will those processes operate?
What kind of rules will the organization have?	Where should management be placed?	How will they be measured?
		How will the organization determine that the processes are working correctly?

Step 2. Define resource requirements

- How much time is needed (per resource)?
- How much money will be spent?
- What personnel will be involved in the change?
- What training will be in place?

Step 3. Identify change agents

- Who are the primary change agents involved?
- Are the primary change agents aware of their responsibilities?

After you gather this data, consider doing an analysis of how your organization and business processes are organized today and how you would like to see them designed in the future. This activity should be led jointly by the OCA team and the cloud or transformation program team.

Finally, when modeling the future state, consider describing changes to the following components of the Cloud Adoption Framework (CAF) People Perspective:

- Culture evolution: Evaluate, incrementally evolve, and codify organizational culture with digital transformation aspirations.
- Transformational leadership: Strengthen leadership capability and mobilize leaders to drive transformational change.

- Cloud fluency: Build digital acumen to confidently and effectively leverage the cloud to accelerate business outcomes.
- Workforce transformation: Enable talent and modernize roles to attract, develop, and retain a digitally fluent and high-performing workforce.
- Change acceleration: Accelerate the adoption of the new ways of working by applying a programmatic change acceleration framework.
- Organization design: Assess and evolve organizational design for alignment with the new cloud ways of working.
- Organizational alignment: Establish ongoing partnerships between organizational structures, business operations, talent, and culture.

FAQ

Q. Who should be involved in this activity?

A. Executive sponsors, project or program leader, change leader, internal service providers (for example, communications, training/learning, and human resources, if they have a role in supporting the change initiative).

Q. What are the inputs to this activity?

A. Business case, discovery phase outputs (MRA and MRP), interviews with executive sponsor and human resources, staffing model, culture assessments, cloud strategy, and business value realization plans.

Q. What are the outputs of this activity?

A. High-level future state organizational model and high-level description of roles and responsibilities.

Q. Why should time be spent on this activity?

A. Investing time to define the desired cloud future state helps align the organization on the destination for the cloud journey.

Q. When do you use it?

A. Use a future state approach to intentionally change the way your company works, and determine how people empower the business strategy. This might result in drastic changes such as

outsourcing, insourcing, or hiring a managed service to deliver aspects of your business. To make these types of decisions around the future state, be sure to involve participants who have diverse experiences from different professions to encourage innovation in the solution space.

Additional steps

To start mobilizing the team and defining the future state:

- Review the business case.
- Review discovery phase outputs.
- Review the cloud strategy and business value realization plans.
- Analyze input from interviews with the executive sponsor, HR, and other stakeholders.
- Participate in operational model design sessions.
- Review staffing models.
- Review any cultural assessments.
- Facilitate creation of the future structure at a high level and ensure buy-in from key stakeholders.
- Work with executive leadership to assess current leadership and determine the key leaders of the future organizational structure.
- Review future state and business requirements.

1.7 Change adoption metrics

Overview

Change adoption metrics are performance measures that monitor and track how the people in your organization are adopting required future state changes in processes, technology use, and ways of working. Metrics might be both qualitative and quantitative, and can include both lagging indicators and leading indicators.

We recommend that you establish an OCA scorecard that tracks both qualitative measures (such as employee perceptions of the change and commitment to change) and quantitative measures (such as percentage of employees who attended scheduled training or heard about the change from their direct manager).

The following guiding principles are critical to the success of cloud adoption and should be measured:

- Leadership is informed and supportive of cloud transformation timeline, milestones, and required organizational support.
- A clear, concise, well-articulated vision of the future and a compelling case for change is understood.
- Stakeholders at all levels have clarity of the change at the personal level. They are aware of what it will take to get there, and they take ownership of the change.
- All employees who are affected by the changes are fully aware, are prepared, and receive timely and relevant enablement and training.
- Program information and support resources are available throughout cloud transformation.

These guiding principles, implemented by a robust culture and change plan, help accelerate business user adoption and program success.

Best practices

In our experience, culture change acceleration metrics are typically lagging indicators instead of leading indicators, as defined in the following table. It is important to track both types of indicators depending on your program initiative goals and objectives.

Measure design principle	Definition	Example measures
Lagging indicator	Measures the success of a change activity (and achievement of a change outcome) after it has happened.	Percentage of staff that agree or strongly agree that training was relevant Percentage attendance at scheduled training
Leading indicator	Measures how the organizat ion is tracking toward achieving a change outcome (such as employees having the skills to perform their roles) at various intervals during the project. Periodica lly measuring change by using lead indicators identifies corrective actions that might be required to ensure that the change outcome is achieved and is sustainable.	Percentage of staff that agree or strongly agree that they have the skills to perform their roles

Metrics typically fall into the four categories listed in the following table. Change acceleration metrics should include both qualitative and quantitative metrics.

Shared vision and strategy	Engagement and alignment of sponsors	Engagement of business users	Skill and competenc y development
• Awareness of program • Messaging effective ness • Alignment	• Commitment • Readiness • Prioritization	• Awareness of resources • Readiness • Understanding of impact	• Training effective ness • Readiness to perform job tasks

Shared vision and strategy	Engagement and alignment of sponsors	Engagement of business users	Skill and competency development
• Impact			

Guidelines

You can use the following data to track metrics (not a complete list):

- Surveys
- Email receipts
- Email link usage
- Evaluations
- Proficiency, metrics
- One-on-one meetings
- Major program events
- Change ambassador feedback

The following table focuses on ways to measure change and people management components.

Change area	What to measure (change outcome or benefit)
Shared vision and strategy	• People understand where the project is headed and their relationship with their group or team. • It's clear how the project will help achieve its vision and goal. • It's clear how the project fits in with other projects. • Changes in practice are clear and meaningful. • The need for change is compelling.

Change area	What to measure (change outcome or benefit)
	• Vision is understood at all levels. • Clear business output and milestones are agreed upon and communicated.
Leader engagement and alignment	• Teams are well supported by leaders. • All key leadership levels communicate the project vision. • Leaders are clear about their roles and accountabilities. • Leaders demonstrate commitment through their actions and behavior. • Leaders model new values and behavior. • Leaders are receptive to new ideas. • Leaders maintain focus when faced with other priorities. • Leaders demonstrate personal commitment. • Leaders provide timely and relevant coaching in new skills, knowledge, and behavior.
People engagement and communications	• Employees are aware of the need for change. • Employees have a sense of urgency for required changes. • There's a clear understanding of benefits for employees. • Communications are delivered to affected stakeholders (planned versus actual). • Stakeholders are identified, fully involved in the program, and listened to. • Employees demonstrate little resistance.

Change area	What to measure (change outcome or benefit)
Performance of project team	• There's clear agreement on a delivery plan. • There's strong team-working and communications across the team. • Common ways of working are understood and demonstrated. • Team responsibilities and accountabilities are clear. • Decision-making doesn't slow change. • Problems are solved quickly.
Skill and competency development	• There's confidence in new skills and knowledge to perform roles in the new environment. • Timely and relevant training is available.
Alignment of processes and organization	• New processes are agreed upon and understood at all levels. • There's agreement on organizational change opportunities and implications. • Actions are taken to align the organization. • Linked projects are well coordinated. • Program and business-as-usual decisions are well linked.
Alignment of people practices and processes (performance and rewards)	• People performance management processes are adapted to enable change. • People development processes are aligned to vision and strategy.

Change area	What to measure (change outcome or benefit)
Addressing culture (behavior and symbol) gaps	• Cultural requirements to sustain desired change are formalized. • Current culture is reviewed and gaps identified. • Actions are taken to close the gaps.
Identification, tracking, and persistence of benefits	• Program goals and objectives are identified. • Changes are embedded and sustained in people's daily routines.

Example of scorecard for OCA change adoption metrics

The following example shows OCA metrics organized by the organizational change acceleration phase or point. Each OCA phase is likely to have a number of desired change outcomes and will therefore require a number of measures.

ID	Org. Change Acceleration Phase	Category	What are we measuring? (Change outcome / benefit)	Target / Metric	Current Status (As of: DATE) Red: <60 Yellow: 60.1 - 79.9 Green: >80
1	Mobilize the Team	Goal and Directions	Clear migration delivery plan agreed to	Migration Inventory - *Migration Plan in place and agreed upon with specific dates and points of contact. Red: less than 60% of the application portfolio has a migration date set, Yellow: 60.1 - 79.9% of the application portfolio has a migration date set, Green: 80%+ of the application portfolio has a migration date set*	**96%**
7	Align Leaders	Empower and Support	Teams are well supported by leaders	Organizational Readiness Survey - *My manager/leader takes an active interest in this effort*	90%
8	Align Leaders	Define / Establish CCoE (CCOE)	CCOE team members are clear about their roles & accountabilities	RACI Workshop and Polling - I understand my role and responsibilities as a CCOE member	0%
9	Align Leaders	Deploy and Operate	CCOE RACI is deployed to all of the other IT Teams	RACI Rollout - *Parties that were impacted by the RACI have received an interactive communication/training/meeting to describe how the roles and responsibilities will function*	0%
10	Align Leaders	Define / Establish CCoE (CCOE)	CCOE demonstrating commitment by meeting regularly and holding productive meetings.	RACI Meeting Cadence - *Meeting on a weekly basis, building out structured backlogs, and demonstrating ownership of tasks.*	80%
11	Align Leaders	Prioritization of Time	Maintain focus when faced with other priorities	Modernization Events - *Establishment of a modernization plan. Modernization discovery cadence.*	0%
12	Align Leaders	Prioritize and Own (Build and Operate in the Cloud)	Demonstrate personal commitment	Modernization Events - *Attendance and active participation in Modernization activities by those identified to participate .*	0%
15	Envision the Future	Prioritization of Time	Clarity on how the Cloud Program integrates with other organizational initiatives, related to resource allocation, emphasis and time, internal to CUSTOMER.	Organizational Readiness Survey - I understand the priority of this change in relation to other initiatives within CUSTOMER	95%

FAQ

Q. Who should be involved in this activity?

A. Executive sponsors, cloud program leader, cloud change leader, internal service providers (for example, communications, training/ learning, and human resources, if they have a role in supporting the change initiative).

Q. Why is it valuable?

A. Change performance measures help you measure and track whether people are effectively transitioning through the required changes. In most projects, the technical, financial, and operational aspects of implementation are closely tracked and monitored, but people issues are often ignored or undiagnosed until they become problems. The high failure rate that characterizes project implementation is associated more closely with the inability to manage people through change, rather than operational or financial factors.

Q. When do you use it?

A. You should assess change adoption metrics at each stage of your project to measure the change and make adjustments to OCA strategies. You can use these metrics with any size project, large or small.

Q. What are the inputs to this analysis?

A. Program charter, business case, feedback from project and functional leadership, cloud transformation program metrics, survey tools (if not available, confirm collection), pre-training and post-training data (if not available, confirm collection), access to program leadership (for focus group interviews), and engagement and readiness activity evaluation forms.

Q. What are the outputs of this exercise?

A. Change risk scorecard, recommended mitigation actions, and clear and tangible KPIs to track and assess the success of the program.

Q. Why should time be spent on this activity?

A. In a data-driven society, organizations rely on metrics to evaluate how they're doing in a particular task, initiative, or resource allocation project. Cloud teams whose functions require cooperation and continuous improvement rely heavily on metrics.

Additional steps

1. Identify the change areas and outcomes or benefits.
2. Develop change measures.
3. Identify the data delivery or data gathering method for each measure.
4. Identify targets for each measure and determine when they will be delivered.

1.8 Program budget

Overview

A program budget is the financial plan for a period of the program, such as one year, or the life of the cloud transformation. When it comes to the cloud, you must learn how to manage, optimize, and estimate costs as you run workloads on AWS. That includes following architectural best practices, exploring cost-optimization strategies, and designing patterns that help you architect cost-efficient solutions on AWS. To gain a deeper understanding of these concepts, your finance team should look into AWS Training and Certification courses on cloud financial management.

For the OCA workstream, understanding the costs related to supporting the people and organizational dimensions of cloud adoption is key to controlling and executing tasks and resources, and mitigating risk. Although the budget can vary across OCA projects, studies have suggested that companies invest between 15 and 30 percent of their total project budget on organizational change acceleration.

Best practices

Program budget requirements fall into the following categories:

- OCA team resources (for example, change management, training, communications, technical writers, instructional designers)
- Material development (for example, communications, internal marketing, translations, printed materials)
- Skills and knowledge (for example, specialty training, instructor-led training, game days, workshops, simulations, certifications)
- Travel and events (for example, organizational readiness assessments, local site visits, instructor-led training, promotional events that drive interest and excitement)
- Software (for example, learning management systems, licenses for instructional design, enrollment fees, reporting fees, webinar conferencing tools)
- Hardware (for example, laptop leases or rentals for training)
- Facilities (for example, venue fees for off-site training, conference rooms, projectors, audio/video equipment)

For budget-constrained organizations, many training and events that were traditionally conducted in a live physical environment can also be delivered virtually and asynchronously to contain costs and provide more inclusivity to global team members.

As a best practice, review the program budget periodically depending on the length of the program or transformation, and adjust it for any new requirements or savings. Work with your program finance team to make sure that they understand the value of change acceleration and the associated program budget.

FAQ

Q. Why is it necessary to manage the program budget?

A. Your OCA investment should be directly aligned with the magnitude of the change and the scope of anticipated activities. Understanding the scope gives you better visibility into forecasting and estimating costs.

Consider budgetary requirements for change acceleration, organizational change management, organizational design, culture, communications, and training resources. Also consider expenses related to the development, deployment, and delivery of training and communication materials, software, hardware, and travel-related expenses.

Q. When should you manage the program budget?

A. To support the creation of a robust program budget, you can anticipate and plan most OCA activities in advance, with inputs from cloud-related assessments such as the MRA. However, unplanned activities can surface throughout the cloud adoption effort and will require investigation, assessment, and approval by the leadership team.

Q. What are the inputs to this process?

A. Inputs into the budgetary process include communication and training resource allocations, travel-related expenses, communication and training material costs, and software-related and hardware-related expenses.

Q. What is the output of this process?

A. The output of this process is an aligned and approved initial program budget that includes all change acceleration activities.

Q. Who should be involved in this process?

A. Involve the following people: executive sponsor, cloud project leader, cloud change leader, internal customer change team liaison, and HR.

Additional steps

To begin defining the program budget, take the following steps as applicable:

1. Review all discovery outputs (for example, readiness assessments, diagnostics) and scoping outputs (for example, cloud plans, roadmaps) to estimate the magnitude of change, scope, timelines, and budgetary implications for the OCA workstream.
2. Interview the internal customer change and HR teams to understand the bandwidth of resources available for the program.
3. Evaluate OCA needs and roles.
4. Estimate the baseline OCA resources required to support the cloud program.
5. Review change readiness assessment findings, update change acceleration resources as needed, and complete the resource budget template.
6. Review, approve, and sign off on the change acceleration budget with the leadership team.
7. Periodically review the program budget against future cloud plans and roadmap to anticipate changes to OCA resources.
8. Periodically review the budget to make sure that it remains below or on target.

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 – 2. Align Leaders
- 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
- Leigh Angus, AWS Strategy, PM & Engagement 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	—	September 12, 2024

AWS Prescriptive Guidance glossary

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

Numbers

7 Rs

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

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

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

A

ABAC

See attribute-based access control.

abstracted services

See managed services.

ACID

See atomicity, consistency, isolation, durability.

active-active migration

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

active-passive migration

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

aggregate function

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

AI

See artificial intelligence.

AIOps

See artificial intelligence operations.

anonymization

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

anti-pattern

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

application control

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

application portfolio

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

artificial intelligence (AI)

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

artificial intelligence operations (AIOps)

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

asymmetric encryption

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

atomicity, consistency, isolation, durability (ACID)

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

attribute-based access control (ABAC)

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

authoritative data source

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

Availability Zone

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

AWS Cloud Adoption Framework (AWS CAF)

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

AWS Workload Qualification Framework (AWS WQF)

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

B

bad bot

A bot that is intended to disrupt or cause harm to individuals or organizations.

BCP

See business continuity planning.

behavior graph

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

big-endian system

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

binary classification

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

bloom filter

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

blue/green deployment

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

bot

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

botnet

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

branch

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

break-glass access

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

brownfield strategy

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

buffer cache

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

business capability

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

business continuity planning (BCP)

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

C

CAF

See AWS Cloud Adoption Framework.

canary deployment

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

CCoE

See Cloud Center of Excellence.

CDC

See change data capture.

change data capture (CDC)

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

chaos engineering

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

CI/CD

See continuous integration and continuous delivery.

classification

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

client-side encryption

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

Cloud Center of Excellence (CCoE)

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

cloud computing

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

cloud operating model

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

cloud stages of adoption

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

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

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

CMDB

See configuration management database.

code repository

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

cold cache

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

cold data

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

computer vision (CV)

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

configuration drift

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

configuration management database (CMDB)

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

conformance pack

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

continuous integration and continuous delivery (CI/CD)

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

CV

See computer vision.

D

data at rest

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

data classification

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

data drift

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

data in transit

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

data mesh

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

data minimization

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

data perimeter

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

data preprocessing

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

data provenance

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

data subject

An individual whose data is being collected and processed.

data warehouse

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

database definition language (DDL)

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

database manipulation language (DML)

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

DDL

See database definition language.

deep ensemble

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

deep learning

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

defense-in-depth

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

delegated administrator

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

deployment

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

development environment

See environment.

detective control

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

development value stream mapping (DVSM)

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

digital twin

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

dimension table

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

disaster

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

disaster recovery (DR)

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

DML

See database manipulation language.

domain-driven design

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

DR

See disaster recovery.

drift detection

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

DVSM

See development value stream mapping.

E

EDA

See exploratory data analysis.

EDI

See electronic data interchange.

edge computing

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

electronic data interchange (EDI)

The automated exchange of business documents between organizations. For more information, see What is Electronic Data Interchange.

encryption

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

encryption key

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

endianness

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

endpoint

See service endpoint.

endpoint service

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

enterprise resource planning (ERP)

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

envelope encryption

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

environment

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

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

epic

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

ERP

See enterprise resource planning.

exploratory data analysis (EDA)

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

F

fact table

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

fail fast

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

fault isolation boundary

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

feature branch

See branch.

features

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

feature importance

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

feature transformation

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

few-shot prompting

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

FGAC

See fine-grained access control.

fine-grained access control (FGAC)

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

flash-cut migration

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

FM

See foundation model.

foundation model (FM)

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

G

generative AI

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

geo blocking

See geographic restrictions.

geographic restrictions (geo blocking)

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

Gitflow workflow

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

golden image

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

greenfield strategy

The absence of existing infrastructure in a new environment. When adopting a greenfield strategy for a system architecture, you can select all new technologies without the restriction

of compatibility with existing infrastructure, also known as brownfield. If you are expanding the existing infrastructure, you might blend brownfield and greenfield strategies.

guardrail

A high-level rule that helps govern resources, policies, and compliance across organizational units (OUs). *Preventive guardrails* enforce policies to ensure alignment to compliance standards. They are implemented by using service control policies and IAM permissions boundaries. *Detective guardrails* detect policy violations and compliance issues, and generate alerts for remediation. They are implemented by using AWS Config, AWS Security Hub CSPM, Amazon GuardDuty, AWS Trusted Advisor, Amazon Inspector, and custom AWS Lambda checks.

H

HA

See high availability.

heterogeneous database migration

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

high availability (HA)

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

historian modernization

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

holdout data

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

homogeneous database migration

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

hot data

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

hotfix

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

hypercare period

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

I

IaC

See infrastructure as code.

identity-based policy

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

idle application

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

IIoT

See industrial Internet of Things.

immutable infrastructure

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

inbound (ingress) VPC

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

incremental migration

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

Industry 4.0

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

infrastructure

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

infrastructure as code (IaC)

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

industrial Internet of Things (IIoT)

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

inspection VPC

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

Internet of Things (IoT)

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

interpretability

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

IoT

See Internet of Things.

IT information library (ITIL)

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

IT service management (ITSM)

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

ITIL

See IT information library.

ITSM

See IT service management.

L

label-based access control (LBAC)

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

landing zone

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

large language model (LLM)

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

large migration

A migration of 300 or more servers.

LBAC

See label-based access control.

least privilege

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

lift and shift

See 7 Rs.

little-endian system

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

LLM

See large language model.

lower environments

See environment.

M

machine learning (ML)

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

main branch

See branch.

malware

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

managed services

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

manufacturing execution system (MES)

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

MAP

See Migration Acceleration Program.

mechanism

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

member account

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

MES

See manufacturing execution system.

Message Queuing Telemetry Transport (MQTT)

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

microservice

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

microservices architecture

An approach to building an application with independent components that run each application process as a microservice. These microservices communicate through a well-defined interface by using lightweight APIs. Each microservice in this architecture can be updated, deployed,

and scaled to meet demand for specific functions of an application. For more information, see Implementing microservices on AWS.

Migration Acceleration Program (MAP)

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

migration at scale

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

migration factory

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

migration metadata

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

migration pattern

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

Migration Portfolio Assessment (MPA)

An online tool that provides information for validating the business case for migrating to the AWS Cloud. MPA provides detailed portfolio assessment (server right-sizing, pricing, TCO

comparisons, migration cost analysis) as well as migration planning (application data analysis and data collection, application grouping, migration prioritization, and wave planning). The MPA tool (requires login) is available free of charge to all AWS consultants and APN Partner consultants.

Migration Readiness Assessment (MRA)

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

migration strategy

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

ML

See machine learning.

modernization

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

modernization readiness assessment

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

monolithic applications (monoliths)

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

use a microservices architecture. For more information, see Decomposing monoliths into microservices.

MPA

See Migration Portfolio Assessment.

MQTT

See Message Queuing Telemetry Transport.

multiclass classification

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

mutable infrastructure

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

O

OAC

See origin access control.

OAI

See origin access identity.

OCM

See organizational change management.

offline migration

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

OI

See operations integration.

OLA

See operational-level agreement.

online migration

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

OPC-UA

See Open Process Communications - Unified Architecture.

Open Process Communications - Unified Architecture (OPC-UA)

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

operational-level agreement (OLA)

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

operational readiness review (ORR)

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

operational technology (OT)

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

operations integration (OI)

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

organization trail

A trail that's created by AWS CloudTrail that logs all events for all AWS accounts in an organization in AWS Organizations. This trail is created in each AWS account that's part of the

organization and tracks the activity in each account. For more information, see Creating a trail for an organization in the CloudTrail documentation.

organizational change management (OCM)

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

origin access control (OAC)

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

origin access identity (OAI)

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

ORR

See operational readiness review.

OT

See operational technology.

outbound (egress) VPC

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

P

permissions boundary

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

personally identifiable information (PII)

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

PII

See personally identifiable information.

playbook

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

PLC

See programmable logic controller.

PLM

See product lifecycle management.

policy

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

polyglot persistence

Independently choosing a microservice's data storage technology based on data access patterns and other requirements. If your microservices have the same data storage technology, they can encounter implementation challenges or experience poor performance. Microservices are more

easily implemented and achieve better performance and scalability if they use the data store best adapted to their requirements.

portfolio assessment

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

predicate

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

predicate pushdown

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

preventative control

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

principal

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

privacy by design

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

private hosted zones

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

proactive control

A security control designed to prevent the deployment of noncompliant resources. These controls scan resources before they are provisioned. If the resource is not compliant with the control, then it isn't provisioned. For more information, see the Controls reference guide in the

AWS Control Tower documentation and see Proactive controls in *Implementing security controls on AWS*.

product lifecycle management (PLM)

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

production environment

See environment.

programmable logic controller (PLC)

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

prompt chaining

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

pseudonymization

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

publish/subscribe (pub/sub)

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

Q

query plan

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

query plan regression

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

R

RACI matrix

See responsible, accountable, consulted, informed (RACI).

RAG

See Retrieval Augmented Generation.

ransomware

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

RASCI matrix

See responsible, accountable, consulted, informed (RACI).

RCAC

See row and column access control.

read replica

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

re-architect

See 7 Rs.

recovery point objective (RPO)

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

recovery time objective (RTO)

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

refactor

See 7 Rs.

Region

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

regression

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

rehost

See 7 Rs.

release

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

relocate

See 7 Rs.

replatform

See 7 Rs.

repurchase

See 7 Rs.

resiliency

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

resource-based policy

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

responsible, accountable, consulted, informed (RACI) matrix

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

responsive control

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

retain

See 7 Rs.

retire

See 7 Rs.

Retrieval Augmented Generation (RAG)

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

rotation

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

row and column access control (RCAC)

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

RPO

See recovery point objective.

RTO

See recovery time objective.

runbook

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

S

SAML 2.0

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

SCADA

See supervisory control and data acquisition.

SCP

See service control policy.

secret

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

security by design

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

security control

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

security hardening

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

security information and event management (SIEM) system

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

security response automation

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

server-side encryption

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

service control policy (SCP)

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

service endpoint

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

service-level agreement (SLA)

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

service-level indicator (SLI)

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

service-level objective (SLO)

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

shared responsibility model

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

SIEM

See security information and event management system.

single point of failure (SPOF)

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

SLA

See service-level agreement.

SLI

See service-level indicator.

SLO

See service-level objective.

split-and-seed model

A pattern for scaling and accelerating modernization projects. As new features and product releases are defined, the core team splits up to create new product teams. This helps scale your organization's capabilities and services, improves developer productivity, and supports rapid

innovation. For more information, see Phased approach to modernizing applications in the AWS Cloud.

SPOF

See single point of failure.

star schema

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

strangler fig pattern

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

subnet

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

supervisory control and data acquisition (SCADA)

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

symmetric encryption

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

synthetic testing

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

system prompt

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

T

tags

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

target variable

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

task list

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

test environment

See environment.

training

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

transit gateway

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

trunk-based workflow

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

trusted access

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

tuning

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

two-pizza team

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

U

uncertainty

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

undifferentiated tasks

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

upper environments

See environment.

V

vacuuming

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

version control

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

VPC peering

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

vulnerability

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

W

warm cache

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

warm data

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

window function

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

workload

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

workstream

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

WORM

See write once, read many.

WQF

See AWS Workload Qualification Framework.

write once, read many (WORM)

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

Z

zero-day exploit

An attack, typically malware, that takes advantage of a zero-day vulnerability.

zero-day vulnerability

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

zero-shot prompting

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

zombie application

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