AWS Organizational Change Acceleration (OCA) 6-Point Framework – 6. Make Culture Change Stick

AWS Prescriptive Guidance

AWS Prescriptive Guidance: AWS Organizational Change Acceleration (OCA) 6-Point Framework – 6. Make Culture Change Stick

Table of Contents

AWS Organizational Change Acceleration (OCA) 6-Point Framework – 6. Make Culture Change Stick

Amazon Web Services (contributors)

February 2025 (document history)

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

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

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

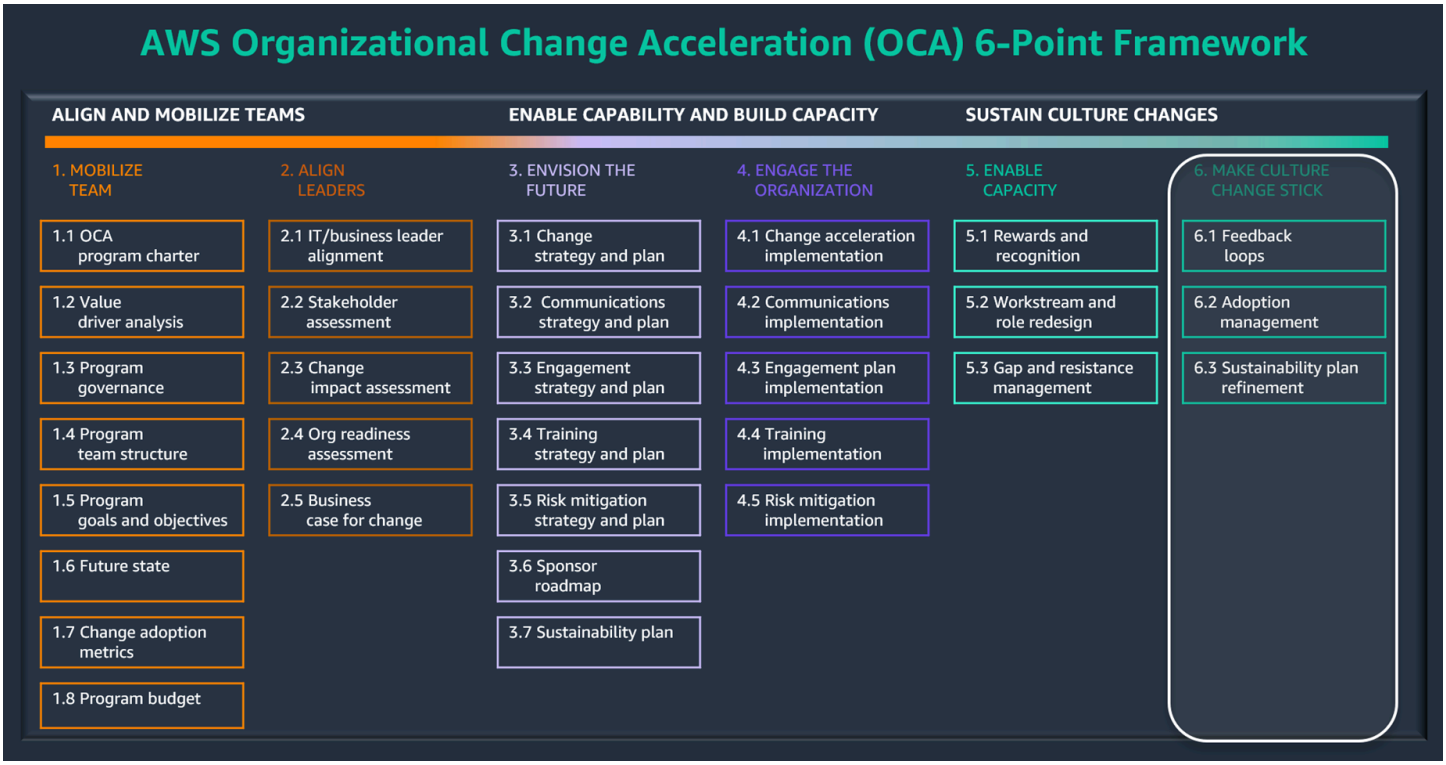

The sixth point, *Make Culture Change Stick*, takes the work of the OCA Framework and builds mechanisms to assess and sustain cloud adoption and culture change over time. In this phase, you create feedback loops for repeatable patterns and lessons learned, actively manage adoption, and create a post-implementation and sustainability plan so that the OCA team can be disbanded, and the changes, behaviors, and culture that have been created can be managed operationally and passively instead of actively. *Make Culture Change Stick* contains three subpoints:

- 6.1 Feedback loops. Establish mechanisms to support two-way information sharing, involve key stakeholders throughout the project, and collect information to monitor communication effectiveness.
- 6.2 Adoption management. Implement the communications strategy and plan to address ongoing communication needs as you implement the cloud strategy.
- 6.3 Sustainability plan refinement. Implement stakeholder-specific plans to address changes that are required to successfully implement the cloud strategy and realize business value from the cloud.

This guide discusses each subpoint of *Make Culture Change Stick* 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 *Make Culture Change Stick* phase of the AWS OCA 6-Point Framework contributes to the following outcomes:

- **Sustainable cultural transformation**: Embedded feedback mechanisms and adoption management processes ensure that cloud-centric practices and mindsets become permanently woven into the organization's foundations and create an enduring culture of innovation.
- **Long-term value realization**: Systematic sustainability planning and continuous feedback loops enable organizations to maintain an accelerated return on investment (ROI) well beyond the initial transformation phase to ensure lasting business value.
- **Institutionalized cloud fluency**: Established mechanisms for knowledge sharing and skill development create self-perpetuating cycles of learning, maintaining, and expanding cloud expertise across the organization over time.
- **Self-sustaining business agility**: Embedded adoption management processes ensure that the organization maintains its ability to rapidly respond to market changes and customer needs without requiring constant, active management.
- **Automated cost optimization**: Established feedback mechanisms help organizations continuously identify and implement cost optimization opportunities and create a self-improving cycle of cloud resource efficiency.
- **Self-reinforcing innovation culture**: Built-in sustainability mechanisms establish innovation practices and enable the ongoing creation of new products, services, and business models without active intervention.
- **Enduring organizational alignment**: Systematic feedback loops maintain continuous alignment between cloud initiatives and business strategy, and ensure lasting strategic coherence.
- **Sustained employee engagement**: Embedded mechanisms for skill development and recognition create self-perpetuating cycles of employee growth and satisfaction, which lead to long-term retention.
- **Consistent market responsiveness**: Institutionalized cloud practices enable the organization to maintain accelerated time-to-market capabilities as a standard operating procedure.

- **Proactive risk management**: Established feedback and monitoring mechanisms enable continuous identification and mitigation of risks, creating a self-sustaining risk management culture.

About the OCA 6-Point Framework guides

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

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

1. Mobilize Team
2. Align Leaders
3. Envision the Future
4. Engage the Organization
5. Enable Capacity
6. Make Culture Change Stick (this guide)

For a comprehensive set of cloud transformation strategies, guidance, and resources, see Accelerating cloud transformation.

6.1 Feedback loops

Overview

Strategic feedback mechanisms are critical drivers of successful cloud transformation that directly impact ROI and business value realization. These systems enable rapid strategy refinement, enhance stakeholder alignment, and create a self-optimizing transformation program.

The implementation framework for feedback loops include the following:

Strategic design

- Establishing clear metrics that are aligned with business objectives
- Defining feedback channels and data collection methods
- Creating rapid response mechanisms for critical insights

Key mechanisms

- Executive interviews and stakeholder forums
- Digital feedback platforms and analytics
- Regular pulse surveys and sentiment analysis
- Cross-functional review sessions

Action protocol

- Real-time insight analysis and prioritization
- Swift implementation of necessary adjustments
- Regular stakeholder communications on actions taken
- Continuous measurement of impact

Success metrics

- Adoption rate acceleration
- Stakeholder satisfaction scores

- Speed of issue resolution
- Resource optimization rates
- Business value realization

By implementing these strategic feedback mechanisms, organizations create a dynamic system that continuously optimizes cloud transformation efforts, maximizes ROI, and ensures sustainable business value creation.

Best practices

Guidelines

To maximize the impact of feedback loops on cloud adoption, consider the following:

- When you implement feedback loops, prioritize areas that have known issues or risks.
- Establish clear response times (for example, within 24 hours for critical feedback).
- Tailor feedback mechanisms to specific stakeholder needs.
- Implement multiple feedback channels to capture diverse perspectives.
- Determine appropriate feedback frequency (for example, weekly for project teams, monthly for executives).
- Ensure data validity through consistent measurement practices.
- Use user-friendly feedback tools to encourage participation.
- Act on feedback promptly and transparently.
- Maintain anonymity when necessary to encourage honest responses.

Opportunity areas

Focus feedback collection on these key areas to drive cloud adoption:

- Project team performance and agility
- Leadership alignment with cloud strategies
- Subject matter expert (SME) and cloud champion engagement
- New employee onboarding effectiveness

- Program evolution and adaptability
- Communication effectiveness across all channels
- Quality and impact of cloud awareness initiatives
- Training material efficacy and skills development
- Timing and relevance of cloud-related events
- Measurement and communication of cloud adoption metrics and business outcomes

Feedback mechanisms

Implement a diverse set of feedback mechanisms to capture comprehensive insights. These might include:

- One-on-one interviews with key stakeholders
- Workstream-specific focus groups
- Regular team, staff, and department meetings
- Engagement and organizational readiness surveys
- Interactive communication portals
- Dedicated project or program feedback channels
- Social media response analysis (for example, tracking cloud-related content engagement)

Integrate feedback collection into your measurement strategy with clear response timelines and cadence. Align feedback cycles with your program's rhythm—for example, collect insights weekly, biweekly, and monthly to inform regular retrospective meetings within your scrum framework. These structured feedback loops not only enhance current initiatives but also create valuable insights that you can use to scale future organizational transformations more effectively. By establishing these systematic feedback mechanisms, you create a continuous learning engine that drives both immediate improvements and long-term organizational growth.

Additional steps

To begin establishing a feedback loop, follow these steps:

1. Determine the method for capturing and sharing feedback.
2. Develop the feedback process. Involve the internal change and HR teams.

3. Develop the feedback tool (surveys, evaluation forms, interview questionnaires, project mailboxes, and so on).
4. Implement and integrate the feedback tool in all change engagement tactics and engagement touch points.
5. Document and report feedback.
6. Demonstrate the impact of feedback by regularly updating stakeholders on how their input has shaped program decisions and actions. This transparency builds trust and encourages continued engagement in the feedback process.

6.2 Adoption management

Overview

Adoption management is the key to successful cloud transformation and serves as the critical bridge between strategic intent and operational reality. In the rapidly evolving landscape of cloud technologies, organizations that excel in adoption management gain a significant competitive advantage and accelerate their journey to becoming cloud-native enterprises.

At its core, adoption management ensures that the substantial investments in cloud infrastructure, tools, and processes translate into tangible business outcomes. In the OCA Framework, adoption management goes beyond technology implementation, because it focuses on the human element that ultimately determines the success of any transformation initiative.

Effective adoption management:

- Catalyzes organizational change by systematically addressing resistance and fostering a cloud-first mindset across all levels of the organization.
- Accelerates value realization through rapid uptake of cloud technologies, which leads to faster innovation cycles and improved time to market.
- Mitigates transformation risks by identifying and addressing adoption barriers early, and preventing costly delays and setbacks.
- Enhances workforce capabilities through targeted upskilling and reskilling initiatives, and by creating a cloud-fluent workforce that's ready for the digital future.
- Aligns cloud initiatives with business objectives to ensure that cloud adoption efforts directly contribute to strategic goals and deliver measurable ROI.
- Cultivates a culture of continuous improvement by establishing feedback loops and adaptive learning mechanisms that drive ongoing optimization.

Adoption management assumes that all foundational OCA plans, including change, communication, risk, and training strategies, have been developed and approved for implementation. It builds upon these foundations to create a cohesive, organization-wide approach to cloud transformation.

In the context of AWS Cloud adoption, effective adoption management is the difference between organizations that just use cloud services and organizations that truly harness the transformative

power of the cloud to revolutionize their business models, customer experiences, and operational efficiencies. The following diagram illustrates how AWS services and the OCA Framework help organizations achieve improved innovation, speed to market, and ROI.

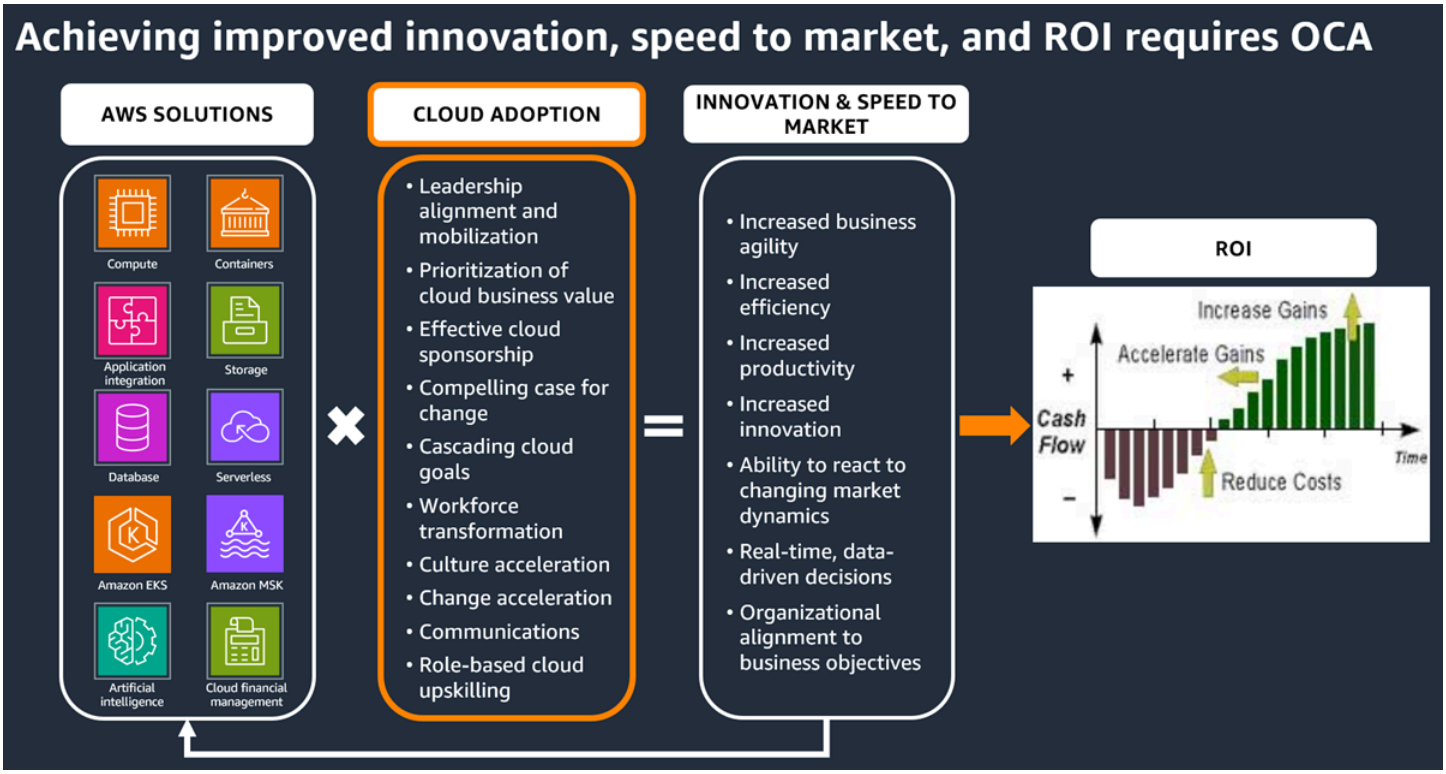

Best practices

To continually monitor adoption and ensure that your organization is on a path to accelerated cloud performance, you can track the indicators of sustained adoption by using a comprehensive checklist. You can include this checklist in your standard agile or project management tool for integrated tracking within your cloud transformation program.

> **Note**
>
> This checklist is representative but not exhaustive. It serves as a starting point for including other items that are specific to your organization's needs.

Example adoption management checklist:

#	Adoption management tasks	Completed	In progress	Not started	Not applicable
1	Leaders are aligned, engaged, and visibly supportive.				
2	Leaders have action plans, tools, and materials to support the initiative.				
3	Managers are involved and participate in transition activities.				
4	SMEs and cloud champions have been onboarded and actively support the initiative as needed.				
5	Key stakeholders are aware of the initiative,				

#	Adoption management tasks	Completed	In progress	Not started	Not applicable
	its benefits, and business value.				
6	Key stakeholders understand and can explain the impacts and changes to their organization.				
7	Internal and external audiences are accounted for in the stakeholder matrix.				
8	Two-way communication vehicles are in place for sharing information.				

#	Adoption management tasks	Completed	In progress	Not started	Not applicable
9	The portal (user information site) is in place and includes the project information vision, overview, impacts, timeline, and FAQs.				
10	The culture change roadmap activities are planned, scheduled, and on track.				
11	The process for measuring and tracking engagement activities is in place.				

#	Adoption management tasks	Completed	In progress	Not started	Not applicable
12	The process for identifying and mitigating risk is in place.				
13	The process for tracking AWS usage and the utilization of new services is in place.				
14	The performance management processes have been updated based on goals, target behaviors, and incentives to reflect cloud priorities.				

#	Adoption management tasks	Completed	In progress	Not started	Not applicable
15	The OCA stakeholder engagement and preparation activities are planned, scheduled, and on track.				
16	Feedback is captured and incorporated into communication and training materials.				
17	Survey results that track readiness activities are trending positively in high percentiles.				

#	Adoption management tasks	Completed	In progress	Not started	Not applicable
18	Training audiences have been identified and mapped to future state training.				
19	The training calendar, curriculum, and course materials have been developed.				
20	Training audiences are enrolled in training.				

#	Adoption management tasks	Completed	In progress	Not started	Not applicable
21	A process is in place for periodically reviewing and modifying training plans to account for new employees, the growing needs of current employees, and new AWS technologies.				
22	Survey results that track training activities are trending positively in high percentiles.				

#	Adoption management tasks	Completed	In progress	Not started	Not applicable
23	Training attendance and completion are within 80-100 percent.				
24	Overall readiness, preparation, and satisfaction survey are tracking within 80-100 percent.				
25	Continuous monitoring and tracking processes are in place.				

Reviewing the adoption management checklist on a quarterly basis provides a mid-range planning horizon for the cloud transformation team and enough flexibility to make adjustments as required.

Additional steps

If cloud adoption is lagging or impeding progress, a systematic approach is required to diagnose and address the issues. Follow the expanded steps in this section to overcome adoption challenges.

1. Conduct a comprehensive adoption audit:
   - Review the adoption management checklist thoroughly, and assess the quality and completeness of each item.
   - Gather quantitative data on cloud usage, application migration progress, and skill development metrics.
   - Conduct qualitative interviews across all organizational levels to uncover hidden barriers and resistances.
2. Perform an organizational pyramid diagnosis. Start at the top and work down the layers, addressing issues at each level.
   a. Executive leadership:
      - If leaders are not fully committed, revisit and strengthen the business case for change.
      - Organize executive cloud immersion sessions to deepen their understanding of cloud benefits.
      - Develop a cloud value dashboard that links cloud initiatives to key business outcomes.
   b. Middle management:
      - Implement a cloud leadership mentoring program that pairs managers with executives who have cloud expertise.
      - Create a cloud champions program to recognize and empower managers who have embraced the cloud.
      - Adjust performance metrics and incentives to align with cloud adoption goals.
   c. Technical teams:
      - Conduct a skills gap analysis to identify specific areas where teams need development.
      - Implement a multi-faceted learning and development plan, including:
        - Role-based AWS certification tracks
        - Hands-on labs and sandboxed environments for experimentation
        - Peer learning groups and internal tech talks
        - Game-based learning challenges to boost engagement
   d. Non-technical staff:
      - Develop introductory courses to build basic cloud literacy across the organization.
      - Create use case showcases that demonstrate how the cloud enables better business outcomes in various departments.

3. Enhance communications and visibility:
   - Establish a cloud adoption portal with real-time dashboards that show adoption progress, success stories, and upcoming initiatives.
   - Implement a regular cadence of town halls and Q&A sessions to address concerns and highlight wins.
   - Develop a storytelling campaign that features employees at all levels who can share their cloud transformation experiences.
4. Optimize the Cloud Center of Excellence (CCoE):
   - Review the composition and mandate of the CCoE to ensure that it's empowered to drive adoption.
   - Implement cross-functional adoption teams within the CCoE to tackle specific adoption challenges.
   - Establish a clear escalation path for adoption blockers that the CCoE can't resolve.
5. Take advantage of AWS resources:
   - Engage AWS account teams and solution architects for tailored adoption acceleration strategies.
   - Use AWS adoption frameworks and tools such as the AWS Cloud Adoption Framework and the AWS Well-Architected Framework.
   - Participate in AWS customer programs such as AWS Executive Insights for senior leader engagement.
6. Implement continuous feedback and iteration:
   - Set up automated sentiment analysis of internal communications and support tickets related to cloud adoption.
   - Conduct quarterly adoption retrospectives to celebrate successes and candidly address setbacks.
   - Establish an adoption innovation fund for employees to propose and implement ideas to accelerate adoption.
7. Address cultural barriers:
   - Work with HR to align hiring, onboarding, and career progression practices with cloud-focused principles.
   - Implement small, targeted interventions to shift behavior toward cloud adoption.
   - Consider organizational structure changes to break down silos that impede cloud adoption.

By systematically working through these steps, organizations can overcome adoption hurdles, accelerate their cloud transformation projects, and fully realize the benefits of their AWS investments. Adoption management is an ongoing process that requires continuous attention and refinement as the organization evolves in its cloud maturity.

6.3 Sustainability plan refinement

Overview

The continued refinement of the sustainability plan (Envision the Future, 3.7) is crucial for ensuring the long-term success of cloud transformation, even after the dedicated transformation team disbands. This plan establishes enduring mechanisms to maintain and evolve cloud adoption practices, and embeds these into the organization's foundations. By focusing on sustainability, organizations can:

- Secure lasting returns on their cloud investments.
- Maintain momentum in innovation and efficiency gains.
- Adapt more readily to emerging cloud technologies and practices.
- Foster a culture of continuous improvement and cloud fluency.

To embed change and ensure sustainability, organizations should:

- Formalize change measurement through quarterly reporting to senior executives.
- Integrate cloud adoption metrics into employee performance plans.
- Allocate dedicated time for employees to monitor and drive cloud adoption.
- Align cloud-related activities with formal processes such as annual performance reviews and compliance training.

Best practices

Sustaining improvement gains over time requires a proactive and systematic approach for creating internal sustainability and ownership. Review the following recommended steps to plan for ownership of the future phases of your organizational acceleration strategy.

1. Identify ongoing OCA needs

When you complete the initial stages of cloud transformation, additional changes might emerge. For example, changing the culture or behavior of one business area might require process changes

in another business area. Or cloud success in some areas of the business might potentially scale to other business units. To determine ownership for future phases of change:

- Review your change planning materials and feedback. What key risks have been identified throughout the process that could be addressed with future change projects? What feedback has emerged that you were not expecting? Where do you see natural next steps for the organization?
- Prioritize the potential future changes. Which changes are essential, and which are useful but not critical? How easy would it be to implement these future change initiatives? Which changes can be implemented with the least amount of effort? Which changes will have the greatest impact on the organization?
- Conduct a high-level change impact assessment of future changes to determine the rough magnitude and scope of the changes.
- Identify future change sponsors. Work with your current change sponsor or senior business leaders to identify the senior executive who has the ability to sponsor a change project and determine the potential business case. To provide approval to move forward with the new project, a senior business leader must be able to see business value to their business unit. Repeat the change process from the definition phase for the new change project.
- Document OCA processes and procedures.
- For future change initiatives, repeat the change framework conducted in this change project, starting with project kick-off and assembling your change team.

2. Transition ongoing ownership

Identify OCA activities that need to be completed or should continue beyond the life of the initial cloud projects. What formal structures and responsibilities do you need to establish or assign to standard roles? How do you plan to transition and reach agreement on ongoing ownership of the changes?

Every change project involves a number of stakeholders and requires a number of people to implement the change plan. Identify the people who should be involved in the ongoing sustainability of the change project.

Set time frames to re-evaluate and review ongoing ownership at intervals after the official completion of the project (for example, every 3 months, 6 months, or 12 months, depending on the project duration).

Here are some potential considerations for ensuring that the change is embedded and sustainable:

- Formalize the change measurement through quarterly reporting to the change sponsor or other senior executive.
- Embed the change into the employee's performance plan or job responsibilities.
- Dedicate a percentage of the employee's weekly time to monitor the change.
- Align the change-related activities with other formal processes and policies (such as the annual performance plan, annual compliance training, hiring plans, and budgeting processes).

Depending on the project's needs, you might need to develop a continuous improvement plan and ongoing review plan to monitor and evaluate your change. A continuous improvement plan might include these sections:

- Purpose
- Governance structure
- Roles and responsibilities
- Calendar of events, including review and feedback sessions
- Ongoing measures of success

3. Transition communications

The objective of this step is to shift the complete management of the communications strategy to the customer organization or sustaining function within the company to ensure that planned communications continue. This might include the following tasks:

- Develop the transition plan.
- Identify ongoing communication champions.
- Meet with communication champions to outline roles, responsibilities, hierarchy, and action items.
- Conduct a transition meeting to confirm next steps.
- Develop the communications roadmap.

Key considerations:

- Make sure that communication champions are ready to accept the transition; that is, the team is adequately trained, has the time to accomplish the tasks, and is technologically ready to support the program.
- Identify the right resources to manage ongoing communications at the project's inception.

4. Transition training

The objective of this step is to provide the organization with a follow-up training plan and to develop additional materials to address training gaps that were discovered after the cloud migration. In addition, transition training involves providing the organization with a post-project archive. Key actions to consider:

- Review feedback, lessons learned, and changes.
- Solicit feedback from participants through the evaluation form.
- Develop additional training support materials as appropriate.
- Document a future state training plan that might include training for new hires, quarterly or event-driven training refreshers, training on new AWS solutions and services, AWS certification and re-certification plans, and so on.
- Create an archive of all associated training materials, such as the project plan, audit and assessment findings, training strategy, curriculum outlines, finalized documentation, evaluation forms, and so on.
- Develop a plan for updating materials.

5. Transition change acceleration metrics

An important factor in sustaining change and cloud adoption is the ongoing collection and monitoring of key metrics. Identify the key metrics at multiple organizational levels that are indicative of achieving the planned business outcomes. These key metrics identify any risks from deviating from your cloud goals. Monitor these metrics frequently at three levels to detect deviations that could impact the attainment of desired business outcomes:

- Organizational change acceleration
- Cloud program, project, and workstream
- Business outcomes

Use the following components of the OCA 6-Point Framework:

- Mobilize Team - 1.5 Project goals and objectives
- Mobilize Team - 1.7 Change adoption metrics
- Align Leaders - 2.5 Business case for change

For each level of metrics, make sure that the following is in place to ensure sustainability:

- Data collection and measurement plan: What are the measures, operational definition, data source, data collection method, and frequency of data collection?
- Plan for automation: How can data collection be automated?
- Responsible, accountable, consulted, informed (RACI) matrix: What are the roles and responsibilities around the monitoring of key metrics?
- Response plan: If a deviation is detected, what is the plan to mitigate and improve each key metric?

6. Obtain leadership sign-off

Determine who needs to approve the completion of change acceleration activities. Organize meetings with the change sponsor and business leaders to discuss the following:

- Any new change projects that might have been raised during the project or during the project review session
- Transition of ongoing ownership (RACI matrix)
- Any outstanding items that need to be addressed before the project can be officially completed
- Approval of project deliverables
- Approval for the project

The following table illustrates an example sign-off record sheet.

Date	Document #	Deliverable	Author	Approved by
Insert date	Insert document reference	Title of change-related deliverable (for example,	People who developed the deliverable	Signature of change sponsor or business

Date	Document #	Deliverable	Author	Approved by
	number (if relevant)	communication strategy and plan, change management strategy and plan, continuous improvement plan)		leader who should approve the deliverable

Additional steps

To ensure the long-term success of your cloud transformation, implement the advanced sustainability measures documented in this section.

1. Implement knowledge transfer programs:
   - Establish a shadowing or reverse shadowing process for complex OCA activities before the cloud transformation team disbands.
   - Create a cloud transformation playbook that documents best practices, lessons learned, and key processes.
   - Develop a mentorship program that pairs cloud-savvy employees with employees who are still developing their skills.
2. Conduct a comprehensive cultural assessment by running a post-implementation survey of cultural characteristics that are crucial to cloud transformation success. Use or modify the following survey questions and apply a Likert scale to measure the results (for example, you can use a scale of 1 to 5: strongly disagree, disagree, neutral, agree, strongly agree).

#	Sample statements	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
1	New ways of working (in the					

#	Sample statements	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
	cloud) are ingrained in our daily operations.					
2	Leadership consistently promotes and exemplifies cloud-first attitudes.					
3	Leadership takes the time to explain why our past ways of working no longer suit our future goals.					

#	Sample statements	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
4	New practices that result from the change effort are superior to old norms.					
5	People who behave and perform in ways that support our new vision are promoted.					

#	Sample statements	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
6	Leadership succession is carefully planned. Executives with traditional mentality will not assume key leadership positions.					
7	New, forward-thinking leaders have been hired.					

#	Sample statements	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
8	Our organization is careful in who we hire. New people are not brought on board if they exhibit traits of a culture that we are trying to move away from.					
9	Leadership (those above my direct manager or supervisor) exhibits new behaviors.					

#	Sample statements	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
10	My managers and supervisors exhibit new behaviors.					
11	My peers exhibit new behaviors.					
12	We are consistently rewarded for behavior that suits the new way of doing things.					

#	Sample statements	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
13	We consistently reinforce the vision related to cloud transformation.					
14	We have formed a new culture that values adaptation to change.					
15	We see new behavior becoming part of the norm.					

#	Sample statements	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
16	Employees across all levels understand how cloud technologies benefit their specific roles.					
17	Our organization rapidly adapts to new AWS services and features.					
18	Cloud proficiency is a key factor in our hiring and promotion decisions.					

#	Sample statements	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
19	Cross-functional collaboration has improved due to our cloud-first approach.					
20	Employees feel empowered to experiment and innovate by using cloud technologies.					

#	Sample statements	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
21	Our organizational structure effectively supports ongoing cloud adoption and innovation.					

3. Establish a cloud governance framework:
   - Develop clear policies and guidelines for cloud usage, security, and cost management.
   - Implement automated compliance checks and remediation processes.
   - Create a cloud financial management strategy to optimize spending and to demonstrate ongoing ROI.
4. Integrate cloud metrics into business key performance indicators (KPIs):
   - Align cloud adoption metrics with overall business performance indicators.
   - Incorporate cloud-related goals into executive and management scorecards.
   - Develop a dashboard that ties cloud usage to business outcomes for leadership visibility.
5. Implement continuous learning mechanisms:
   - Establish a regular cadence of internal tech talks and knowledge-sharing sessions.
   - Create a cloud innovation lab where employees can experiment with new AWS services.
   - Develop a curriculum for ongoing cloud education, including both technical and non-technical tracks.
6. Refine OCA processes:
   - Develop a change advisory board specifically for cloud-related initiatives.
   - Implement a formal process for evaluating and adopting new AWS services.

- Create change management templates that are tailored for cloud-specific projects.

7. Conduct regular sustainability audits:
   - Perform bi-annual reviews of the sustainability plan, adjusting as needed based on organizational changes and new AWS capabilities.
   - Engage third-party experts to provide an external perspective on your cloud sustainability efforts.
   - Benchmark your cloud sustainability practices against industry leaders and AWS case studies.
8. Evolve the CCoE:
   - Transition the CCoE from a project-focused entity to a strategic driver of ongoing cloud innovation.
   - Rotate membership in the CCoE to bring in fresh perspectives and spread cloud expertise.
   - Empower the CCoE to drive cross-functional cloud initiatives and remove adoption barriers.
9. Develop a long-term cloud talent strategy:
   - Create cloud-specific career paths within your organization.
   - Partner with universities and coding bootcamps to develop a pipeline of cloud experts.
   - Implement a cloud ambassador program to evangelize cloud adoption both internally and externally.

A well-crafted and diligently executed sustainability plan is the keystone of long-term cloud transformation success. By systematically addressing ongoing needs, transitioning ownership, and embedding cloud-first practices into the organization's culture and processes, companies can ensure that their cloud investments continue to yield benefits well into the future. Regular evaluation and refinement of the sustainability plan, coupled with strong leadership commitment, will drive continuous improvement and innovation in the organization's cloud journey.

Resources

References

- Accelerating your return on cloud investment by adopting a strategic transformation and change methodology
- AWS Change Acceleration 6-Point Framework and Organizational Change Management Toolkit
- AWS Organizational Change Acceleration (OCA) 6-Point Framework – 1. Mobilize Team
- AWS Organizational Change Acceleration (OCA) 6-Point Framework – 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 Cloud Adoption Framework: People Perspective
- AWS Well-Architected Framework
- AWS Executive Insights

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
- Travis McNeal, AWS Change Acceleration Lead

Document history

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

Change	Description	Date
Initial publication	—	February 28, 2025

AWS Prescriptive Guidance glossary

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

Numbers

7 Rs

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

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

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

A

ABAC

See attribute-based access control.

abstracted services

See managed services.

ACID

See atomicity, consistency, isolation, durability.

active-active migration

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

active-passive migration

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

aggregate function

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

AI

See artificial intelligence.

AIOps

See artificial intelligence operations.

anonymization

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

anti-pattern

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

application control

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

application portfolio

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

artificial intelligence (AI)

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

artificial intelligence operations (AIOps)

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

asymmetric encryption

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

atomicity, consistency, isolation, durability (ACID)

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

attribute-based access control (ABAC)

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

authoritative data source

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

Availability Zone

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

AWS Cloud Adoption Framework (AWS CAF)

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

AWS Workload Qualification Framework (AWS WQF)

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

B

bad bot

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

BCP

See business continuity planning.

behavior graph

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

big-endian system

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

binary classification

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

bloom filter

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

blue/green deployment

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

bot

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

botnet

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

branch

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

break-glass access

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

brownfield strategy

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

buffer cache

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

business capability

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

business continuity planning (BCP)

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

C

CAF

See AWS Cloud Adoption Framework.

canary deployment

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

CCoE

See Cloud Center of Excellence.

CDC

See change data capture.

change data capture (CDC)

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

chaos engineering

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

CI/CD

See continuous integration and continuous delivery.

classification

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

client-side encryption

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

Cloud Center of Excellence (CCoE)

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

cloud computing

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

cloud operating model

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

cloud stages of adoption

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

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

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

CMDB

See configuration management database.

code repository

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

cold cache

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

cold data

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

computer vision (CV)

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

configuration drift

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

configuration management database (CMDB)

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

conformance pack

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

continuous integration and continuous delivery (CI/CD)

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

CV

See computer vision.

D

data at rest

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

data classification

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

data drift

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

data in transit

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

data mesh

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

data minimization

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

data perimeter

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

data preprocessing

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

data provenance

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

data subject

An individual whose data is being collected and processed.

data warehouse

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

database definition language (DDL)

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

database manipulation language (DML)

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

DDL

See database definition language.

deep ensemble

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

deep learning

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

defense-in-depth

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

delegated administrator

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

deployment

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

development environment

See environment.

detective control

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

development value stream mapping (DVSM)

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

digital twin

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

dimension table

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

disaster

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

disaster recovery (DR)

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

DML

See database manipulation language.

domain-driven design

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

DR

See disaster recovery.

drift detection

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

DVSM

See development value stream mapping.

E

EDA

See exploratory data analysis.

EDI

See electronic data interchange.

edge computing

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

electronic data interchange (EDI)

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

encryption

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

encryption key

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

endianness

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

endpoint

See service endpoint.

endpoint service

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

enterprise resource planning (ERP)

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

envelope encryption

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

environment

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

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

epic

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

ERP

See enterprise resource planning.

exploratory data analysis (EDA)

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

F

fact table

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

fail fast

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

fault isolation boundary

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

feature branch

See branch.

features

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

feature importance

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

feature transformation

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

few-shot prompting

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

FGAC

See fine-grained access control.

fine-grained access control (FGAC)

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

flash-cut migration

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

FM

See foundation model.

foundation model (FM)

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

G

generative AI

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

geo blocking

See geographic restrictions.

geographic restrictions (geo blocking)

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

Gitflow workflow

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

golden image

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

greenfield strategy

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

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

guardrail

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

H

HA

See high availability.

heterogeneous database migration

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

high availability (HA)

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

historian modernization

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

holdout data

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

homogeneous database migration

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

hot data

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

hotfix

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

hypercare period

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

I

IaC

See infrastructure as code.

identity-based policy

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

idle application

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

IIoT

See industrial Internet of Things.

immutable infrastructure

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

inbound (ingress) VPC

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

incremental migration

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

Industry 4.0

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

infrastructure

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

infrastructure as code (IaC)

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

industrial Internet of Things (IIoT)

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

inspection VPC

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

Internet of Things (IoT)

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

interpretability

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

IoT

See Internet of Things.

IT information library (ITIL)

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

IT service management (ITSM)

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

ITIL

See IT information library.

ITSM

See IT service management.

L

label-based access control (LBAC)

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

landing zone

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

large language model (LLM)

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

large migration

A migration of 300 or more servers.

LBAC

See label-based access control.

least privilege

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

lift and shift

See 7 Rs.

little-endian system

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

LLM

See large language model.

lower environments

See environment.

M

machine learning (ML)

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

main branch

See branch.

malware

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

managed services

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

manufacturing execution system (MES)

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

MAP

See Migration Acceleration Program.

mechanism

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

member account

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

MES

See manufacturing execution system.

Message Queuing Telemetry Transport (MQTT)

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

microservice

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

microservices architecture

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

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

Migration Acceleration Program (MAP)

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

migration at scale

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

migration factory

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

migration metadata

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

migration pattern

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

Migration Portfolio Assessment (MPA)

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

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

Migration Readiness Assessment (MRA)

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

migration strategy

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

ML

See machine learning.

modernization

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

modernization readiness assessment

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

monolithic applications (monoliths)

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

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

MPA

See Migration Portfolio Assessment.

MQTT

See Message Queuing Telemetry Transport.

multiclass classification

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

mutable infrastructure

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

O

OAC

See origin access control.

OAI

See origin access identity.

OCM

See organizational change management.

offline migration

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

OI

See operations integration.

OLA

See operational-level agreement.

online migration

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

OPC-UA

See Open Process Communications - Unified Architecture.

Open Process Communications - Unified Architecture (OPC-UA)

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

operational-level agreement (OLA)

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

operational readiness review (ORR)

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

operational technology (OT)

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

operations integration (OI)

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

organization trail

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

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

organizational change management (OCM)

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

origin access control (OAC)

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

origin access identity (OAI)

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

ORR

See operational readiness review.

OT

See operational technology.

outbound (egress) VPC

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

P

permissions boundary

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

personally identifiable information (PII)

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

PII

See personally identifiable information.

playbook

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

PLC

See programmable logic controller.

PLM

See product lifecycle management.

policy

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

polyglot persistence

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

best adapted to their requirements. For more information, see Enabling data persistence in microservices.

portfolio assessment

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

predicate

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

predicate pushdown

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

preventative control

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

principal

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

privacy by design

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

private hosted zones

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

proactive control

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

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

product lifecycle management (PLM)

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

production environment

See environment.

programmable logic controller (PLC)

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

prompt chaining

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

pseudonymization

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

publish/subscribe (pub/sub)

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

Q

query plan

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

query plan regression

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

R

RACI matrix

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

RAG

See Retrieval Augmented Generation.

ransomware

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

RASCI matrix

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

RCAC

See row and column access control.

read replica

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

re-architect

See 7 Rs.

recovery point objective (RPO)

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

recovery time objective (RTO)

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

refactor

See 7 Rs.

Region

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

regression

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

rehost

See 7 Rs.

release

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

relocate

See 7 Rs.

replatform

See 7 Rs.

repurchase

See 7 Rs.

resiliency

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

resource-based policy

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

responsible, accountable, consulted, informed (RACI) matrix

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

responsive control

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

retain

See 7 Rs.

retire

See 7 Rs.

Retrieval Augmented Generation (RAG)

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

rotation

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

row and column access control (RCAC)

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

RPO

See recovery point objective.

RTO

See recovery time objective.

runbook

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

S

SAML 2.0

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

SCADA

See supervisory control and data acquisition.

SCP

See service control policy.

secret

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

security by design

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

security control

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

security hardening

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

security information and event management (SIEM) system

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

security response automation

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

server-side encryption

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

service control policy (SCP)

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

service endpoint

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

service-level agreement (SLA)

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

service-level indicator (SLI)

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

service-level objective (SLO)

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

shared responsibility model

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

SIEM

See security information and event management system.

single point of failure (SPOF)

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

SLA

See service-level agreement.

SLI

See service-level indicator.

SLO

See service-level objective.

split-and-seed model

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

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

SPOF

See single point of failure.

star schema

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

strangler fig pattern

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

subnet

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

supervisory control and data acquisition (SCADA)

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

symmetric encryption

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

synthetic testing

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

system prompt

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

T

tags

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

target variable

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

task list

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

test environment

See environment.

training

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

transit gateway

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

trunk-based workflow

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

trusted access

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

tuning

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

two-pizza team

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

U

uncertainty

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

undifferentiated tasks

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

upper environments

See environment.

V

vacuuming

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

version control

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

VPC peering

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

vulnerability

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

W

warm cache

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

warm data

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

window function

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

workload

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

workstream

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

WORM

See write once, read many.

WQF

See AWS Workload Qualification Framework.

write once, read many (WORM)

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

Z

zero-day exploit

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

zero-day vulnerability

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

zero-shot prompting

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

zombie application

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