aws

Strategies for migrating your contact center to Amazon Connect

AWS Prescriptive Guidance

AWS Prescriptive Guidance: Strategies for migrating your contact center to Amazon Connect

Table of Contents

Strategies for migrating your contact center to Amazon Connect

Jag Jhutty, Amazon Web Services (AWS)

December 2024 ([document history](#))

This article defines the objectives and targeted business outcomes of a contact center migration to Amazon Connect. It explains how you can plan the migration, get buy-in from the appropriate stakeholders, carry out the migration, and cut over.

Your contact center is a gateway to your brand and business. Each interaction with an agent, supervisor, or chatbot leaves an impression on your customer. [Amazon Connect](#) is a cloud-based contact center service that enables you to provide personalized customer experiences and deliver exceptional customer service. Amazon Connect provides the following features:

- Omnichannel: Customers can interact with the call center by using a channel of their choosing. You can provide rich digital experiences beyond voice, such as chat, SMS, and social media.
- Consumption-based billing: There are no licenses, contracts, or usage commitments. With Amazon Connect, you pay for what you use.
- Scalability: Amazon Connect is cloud-based, so it scales up and down dynamically to meet demand without your intervention. It automatically handles large call volumes during peak events without requiring you to pay for unused capacity.
- Agility: The frequent release of [new features](#) enables you to stay at the forefront of innovation and customer experiences. New features are ready to activate without requiring any upgrades. Feature roadmaps are customer-driven, based on customer requests, security and reliability points, and operational improvements.
- AI and ML capabilities: You can use built-in artificial intelligence (AI) and machine learning (ML) to personalize and automate interactions, understand customer sentiment, authenticate callers, and enable capabilities such as interactive voice response (IVR) and chatbots.

An independent [Forrester report](#) from June 2020 analyzed six Amazon Connect customers and found that it:

- Reduced total cost of ownership (TCO): 241 percent ROI compared to other contact center providers, reduced subscription and usage costs by 31 percent.

- Deflected and streamlined calls: reduced call volume routing by up to 24 percent.

- Improved visibility: reduced supervisor effort by up to 20 percent due to improvements in reporting and metrics dashboards.

- Simplified management: reduced system administrator efforts by up to 60 percent.

- Boosted customer experience: shortened average handle time (AHT) by up to 15 percent.

- Provided dependability and agility at scale.

This article is for decision makers (for example, the director of infrastructure) who are interested in moving to Amazon Connect because they are unhappy with their existing contact center or they are researching alternatives before an upcoming contract renewal. The article assumes some technical knowledge and familiarity with contact center terminology but no AWS expertise. It provides additional details so that you can forward this article to architects or other technical people within your teams and get their perspective. We also encourage you to discuss the contents of this article with your leadership (for example, corporate executives), recommend a further look at Amazon Connect, and initiate a conversation with your AWS account manager.

Overview

Pillars of successful migration

To carry out a successful contact center migration, you shouldn't view the migration as only a technology delivery project—you should approach it from multiple perspectives. Otherwise, you might overlook vital preparations such as staff training and operating model changes. These non-technology considerations are crucial in ensuring overall success.

The pillars illustrated in the following diagram are perspectives and capabilities described in the AWS Cloud Adoption Framework (AWS CAF). This framework provides best practice guidance to help you digitally transform and accelerate your business outcomes through innovative use of AWS. Each perspective covers a set of capabilities that stakeholders own or manage in the contact center transformation and migration process.
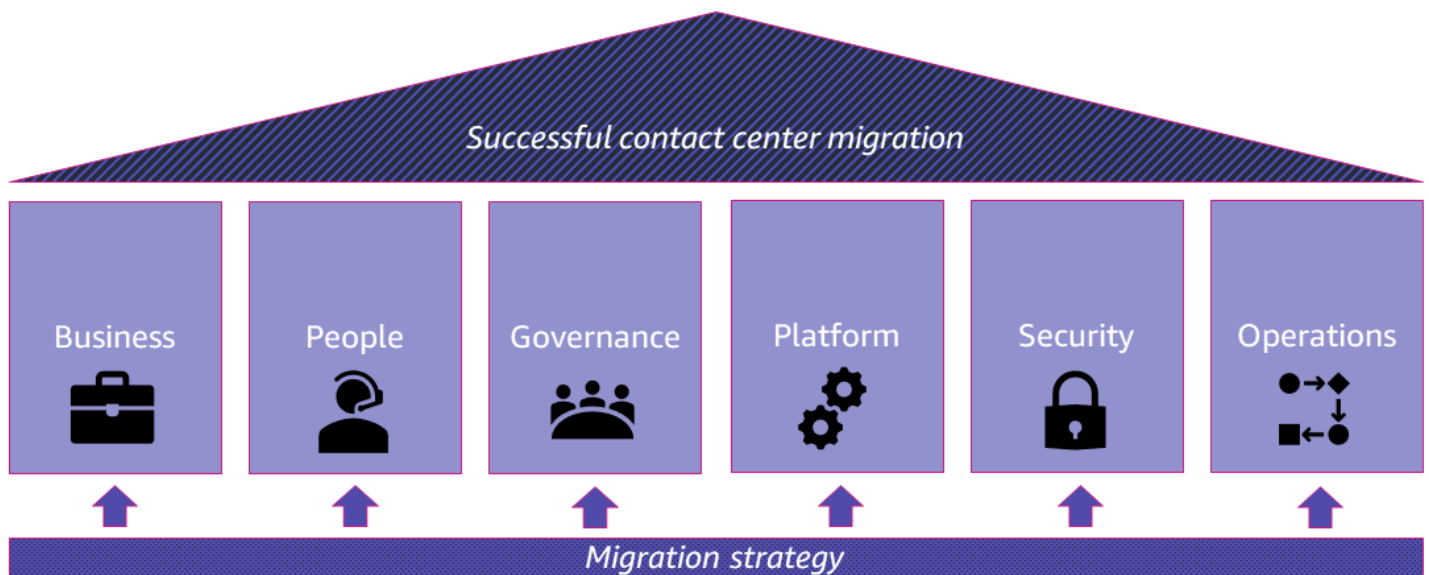
Moving users (customers, agents, and operators) to a new platform and tool set is a considerable amount of work. Contact center migrations require thorough planning, whether you are taking your existing on-premises contact center journeys to the cloud, or refactoring the whole customer and agent experience.

The following sections discuss approaches and best practices to plan, manage, and complete migrations to Amazon Connect.

Primary vision

A successful contact center migration starts with business requirements and then focuses on people, processes, and technology.

Start planning your Amazon Connect migration by first developing a primary vision statement. This should be a general principle that guides the direction of decision-making. You can then define more specific guiding principles for particular decision areas within the bounds of this general principle.

For example, the primary vision statement for your project might answer the question, "What does success look like?" as follows: "*Minimal user disruption* (in order of significance: *customers, agents, system operators*) while migrating service lines at *pace*."

Notice the emphasis on the following phrases:

- *Minimal user disruption* – Depending on your contact center's opening hours and backend systems, it might not be possible to avoid downtime entirely during the migration. Be realistic and consider whether the expected disruption is tolerable compared with the time and effort required to complete the migration without any downtime. Accepting minimal disruption rather than no disruption might reduce risks in other areas of project delivery or provide significant cost savings. For example, you might decide to circulate a new web address to users for accessing the new Amazon Connect desktop instead of migrating an existing web address. This helps avoid the effort and expense of signing new domain certificates and having to manage a web address cutover.

- *User list in order of significance* – Customers, agents, and system operators have different priorities during a migration. Generally, the highest priority is to avoid disruption to your customers, even if it means additional disruption to agents and backend system operators.

- *Pace* – It is costly, both financially and resource-wise, to operate more than one contact center platform during the migration. Your objective should be to keep the dual-system period as short as practical. The longer it is, the greater the cost, the burden on operators, and the risk of human errors such as making changes on the wrong platform. Balance rigor and depth with the need to move rapidly. Develop a realistic delivery plan and try to follow it.

Targeted business outcomes

Keep these business outcomes in mind when you plan your contact center migration:

- **Increased business agility** – Deliver new capabilities into production rapidly and safely. For example, sentiment analysis and big data call transcript crawling help you gather near real-time insight into customer communications and enable you to optimize your products and services based on their needs. After you identify and implement these features, you can deliver them by using DevOps principles, which encourage collaboration among your developers and operators, and use infrastructure as code (IaC) tools and continuous integration and continuous delivery (CI/CD) pipelines to manage builds and automate testing. Avoid repeating steps manually wherever possible to avoid human error, which can introduce bugs into the implementation process.

- **Improved total cost of ownership (TCO), especially in early stages** – Rework costs time and effort. To get key decisions right the first time, allocate sufficient time to the discovery and design phases of migration. Infrastructure decisions are difficult to alter without significant cost, so consult with the appropriate stakeholders. For example, changing the encryption policy for call recordings might require additional infrastructure components, so make sure that your security compliance teams approve the encryption policy before you start implementation. Get sign-off on designs before moving into the build phase.

- **Agile customer experience** – Use agile methodologies to rapidly and iteratively develop caller journeys. Unlike infrastructure components, contact flows and user journeys are easy to alter, so start early with a basic flow and iterate frequently with stakeholders to reach the desired state. It's easy to add a message prompt or to alter menu options in Amazon Connect —no programming knowledge is required. Your objective should be to deliver the right user journey, not to rigidly follow the journey that you originally designed. Iterating frequently gives stakeholders the ability to tweak the journey as it matures and feedback is received.

- **Smooth and timely service introduction** – User training, process changes, and service desk changes are often overlooked until the project is close to completion. The new contact center has to be accepted into your organization's business as usual (BAU) operations as well as meeting the go-live date. Without a proper handover, the project team will not be able to recede and BAU teams will not be prepared to use the new platform. Make your project's integration into BAU operations a gate for go-live approval. It is vital to agree on platform ownership before you go live. Engage service introduction and operating model stakeholders from the beginning of the project, and keep them engaged throughout.

- **Introduce new, differentiating capabilities to improve customer satisfaction (CSAT) scores** – Ask yourself whether the user experience can be simplified or improved by Amazon Connect. Don't limit yourself to lifting and shifting your current call center to the cloud. Use Amazon Connect features to improve the user (customer and agent) experience or to simplify the technical implementation of your platform. With relatively little effort, you can incorporate new

Amazon Connect capabilities into your call center and see a significant improvement in your CSAT scores.

Agile methods to accelerate delivery and innovation

We recommend that you use an agile methodology in conjunction with DevOps and CI/CD practices as the underpinnings of your migration to Amazon Connect. These practices become the foundation for a dynamic, user-focused, and experiment-driven approach to customer experience.

If you have a compelling business reason to initially migrate your contact center as is to Amazon Connect, without adding new features, we still highly recommend adapting an agile approach to enable experimentation and the continuous improvement of the customer experience over time.

Borrowing from [Amazon's business approach to transformation](#), we recommend a *think big, start small, go fast* approach. You start by clarifying your business goals and focus areas, and brainstorm with key stakeholders to define and align on key opportunities for innovation. You then work back from the customer to understand who they are, what they need, and how to improve their experience. From there, you define and prioritize key initiatives to create a minimum lovable product (MLP) that drives business outcomes and provides immediate impact within the initial agile sprint. Establishing an Amazon Connect technical foundation and the agile delivery framework during the initial sprint builds the foundation of the customer experience (CX) flywheel, which is depicted in the following diagram.

CX flywheel

Subsequent sprints are prioritized by working back from customer needs, and organized by additional functionality, additional users and business units, or a combination of the two. The following diagram shows a typical agile sprint process. Change management (CM) activities underpin the agile sprint process and ensure that the organization is keeping pace with technology delivery.

Connect agile delivery with organizational change management (CM)

After teams and stakeholders agree on a multi-phase migration and transformation plan (as discussed in the following sections), the initial agile sprint establishes the foundation of an Amazon Connect contact center, which provides a common baseline of capability, prepares the flywheel mechanism for accelerating transformation, and defines the mechanisms for continuous improvement. The key elements of this foundation include:

- Deploying Amazon Connect on a secure, high-performing, resilient, and efficient AWS infrastructure.

- Configuring contact flows that define the customer experience and establishing design conventions for consistent experiences.

- Developing representative experiences such as customer identification and lookups.

- Setting up the business administration console.

- Integrating critical third-party systems.

- Configuring the data model and data pipeline, such as how to access Amazon Connect data from a data lake or data warehouse.

- Creating a DevOps operational runbook.

These elements are the building blocks for delivering operational fundamentals with next-generation capabilities to elevate the customer experience and reduce operational costs. They are the first items to be used by a project, so they should be prioritized. The foundation is the catalyst for additional sprints and becomes the enabler for continuous experimentation and improvement.

Project phases and workstreams

In the context of a contact center migration project, *sprint*, *workstream*, and *phase* have the following meanings:

- A *sprint* is a time-bound collection of activities that are delivered by different workstreams. For example, each sprint could be two weeks long.

- A *workstream* is a team-bound collection of activities associated with a set of technology components or scope. Sprints include workstream activities. For example, AWS account and landing zone creation can be included in a technical foundation workstream, which involves architect and developer team resources. Mapping customer experiences and recording call prompts should be handled by a different, user journey-related workstream, because these tasks involve business stakeholders and service line owners.

- A *phase* is a goal-oriented collection of activities across workstreams. Phases usually end at milestones, and reaching these milestones means that the project progresses to the next phase. For example, the design phase involves creating documents that are appropriate to each workstream, such as architectural diagrams, build specifications, and high-level design documents. The design phase is completed when these documents are approved by the necessary stakeholders.

Well-defined and autonomous workstreams improve overall project agility. Basing workstreams on specific teams and roles gives team members autonomy in prioritizing sprint backlog items. It also creates boundaries between workstreams, so you can identify and track dependencies, and provides clear accountability.

The high-level plan in the following diagram shows the parallel workstreams and sequence of typical activities in an example contact center migration project.
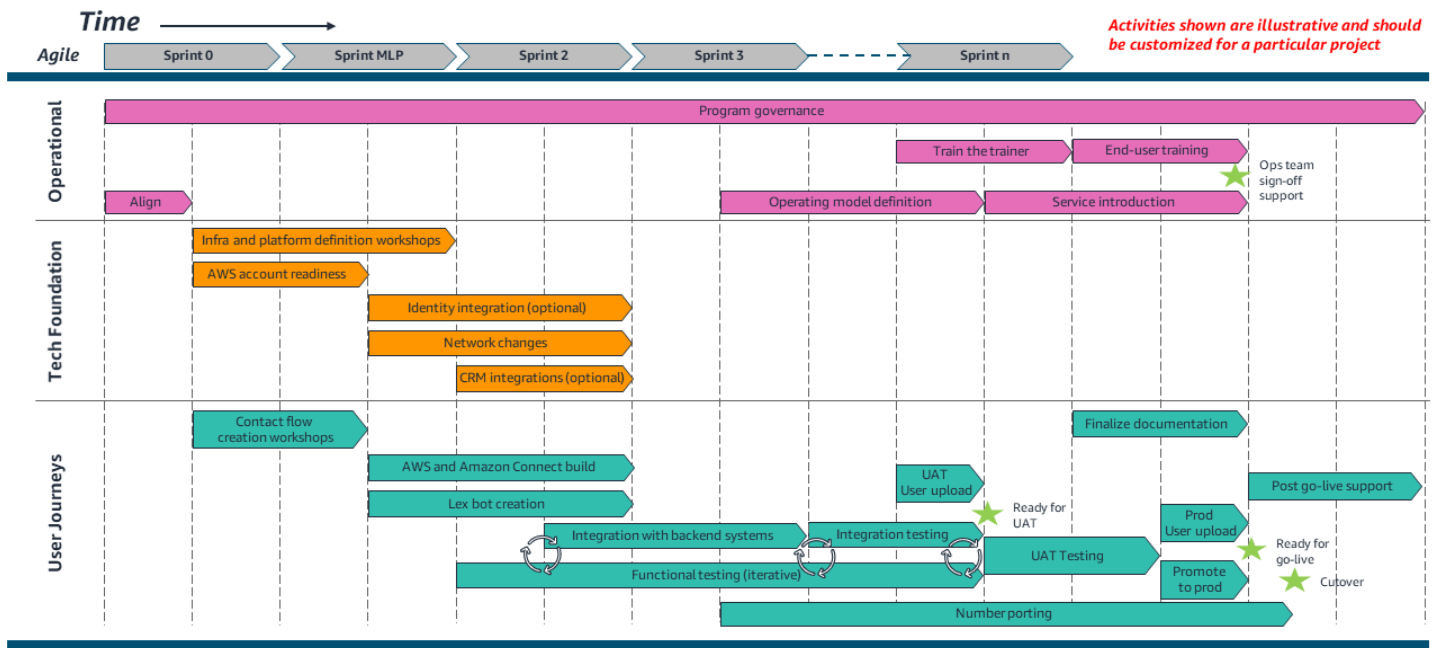
We recommend that you run at least three parallel workstreams: *operational*, *technical foundation*, and *user journeys*. The phasing and approach to project activities differ, depending on the nature of the workstream. Each workstream requires a different delivery approach, as explained in the following sections. As the diagram illustrates:

- Tasks within each workstream are bundled into agile sprints.

- Sprint 0 is a collection of early tasks focused on project kick-off, discovery, planning, and design.

- Sprint MLP is a collection of activities for creating a minimum lovable product (MLP) that future sprints can iterate on to provide end-state target capabilities. For example, the MLP could deliver a relatively straightforward caller journey to a small group of agents. After the platform is live and proven stable for the MLP use cases, future sprints (Sprints 2, 3, and so on in the diagram) can iterate rapidly to deliver innovative capabilities.

- Each project and environment is different, so the diagram doesn't provide specific timelines. Use this plan as a starting point for discussions with stakeholders during the initial project planning phase. Determine which activities are relevant, identify any activities that should be added, and determine their estimated duration.

Operational workstream

The operational workstream supports the technical foundation and user journey workstreams. The majority of non-technical activities that are crucial to overall migration success are part of this workstream.

This workstream involves decisions that can be changed or reversed with little effort or impact. Product specifications that are based on how people work and engage are rarely right the first time—there are many stakeholders and voices to consider. It is important to engage early and consult widely and often, so an agile and iterative approach makes sense for this workstream. You start with an early draft of an operating model or training materials and iterate frequently and rapidly to get to the final product.

The operational workstream consists of five phases: project governance, alignment, operating model definition, service introduction, and training.

Program governance

Program governance activities run throughout the entire timeline of a migration project. Regardless of the stage the project is in, activities should be regular (scheduled recurring meetings), transparent (project team is given the opportunity to raise risks and issues frankly), and with engaged governance (leaders are empowered and willing to make decisions or escalate accordingly). These are vital for highlighting and resolving issues promptly and effectively.

Alignment

This is the first formal activity of the project and focuses on aligning the project scope to business outcomes. Alignment provides an opportunity to validate and adjust earlier plans and estimates based on discussions with stakeholders.

Key actions during this activity include:

- Discover high-level customer use cases, current technical and business pain points, and opportunities for improvement.

- Discuss and agree on desired business outcomes, determine their relative priorities, and identify success criteria.

- Develop a high-level solution design, which is used to define scope and technology choices in this early phase. This high-level design provides direction to accelerate low-level design activities in later phases.

- Validate timelines and implementation costs.

Operating model definition

The activities in this phase define *who* will use the contact center solution and *how* the solution will be managed. The operating model is not a procedural document, a runbook, or a configuration file. For example, it shouldn't explain how to pull logs and attach them to a support ticket, or provide screenshots of this procedure. Instead, it should identify who should pull the logs and which queue or vendor they should be sent to.

The operating model definition should include:

- A responsible, accountable, supported, consulted, and informed (RASCI) matrix, so each team understands its roles and responsibilities, and how they will interact with other teams. The following provides an excerpt from a RASCI matrix.

RACI Defined: R – Who is responsible for doing the actual work for the task A – Who is accountable for the success of the task and is the decision-maker S – Who provides support during the implementation of the activity / process / service C – Who needs to be consulted for details and additional info on requirements I – Who needs to be kept informed of major updates	Process Activity	Business					Amazon Connect CoE							AWS Platform CoE			Salesforce CoE		Notes
		Overall CX Lead	Service Line CX Owner	Governance	Security	Business Analyst	Contact Center Product Owner	Amazon Connect Architect	Amazon Connect Engineer	DevOps Engineer	Contact Center Operations	Telecoms Engineer	Data Analyst	AWS Platform Owner	AWS Architect	DevOps Engineer	SF Platform Owner	SF Admin	
Cloud Architecture	Cloud Architecture Design						C	C						A	R	S			
	Design Infrastructure to support contact flows		S				A	R	C	I					S				
	S3 Lifecycle-Definition			A	C		I	R											
	Terraform IaC & Pipeline (For Contact Center Design & Tasks)		I				I	A	C	R					C	S			
	Github IaC & Pipeline (For Contact Center Design & Tasks)		I				I	A	C	R					C	S			
	KMS Customer Managed Key (CMK) Rotation	I	I	I	A		C	C		R						I			
Amazon Connect Operations																			
	User MACD (Moves, Additions, Changes, Deletions)		A							R									
	User Hierarchies Management		A		C					I	I	R							
	Phone Number Management eg. Claiming & Releasing Numbers						A	I	R										
	Queues - Definition		A		C		R	C											
	Queues - Implementation						A	R											

- Process flow swimlanes that define the end-to-end activities and who is responsible for each activity. For example, there should be a process flow for engaging out-of-hours support so it's clear who gets paged, what happens if they can't be reached, who logs the support ticket, and how the business criticality is judged. Another example is the emergency queuing message. The process flow should show who decides that it needs to be initiated, and what data they should use to make that decision.

The operating model is typically defined in the second half of a project, because you have to finalize the solution design and user journeys before you can define the processes to manage them

accurately. However, we recommend that you line up stakeholders early in the process and reserve their time for the later phases of the project.

Gather examples of similar documents from your organization that you can use as templates. This will facilitate review and sign-off by stakeholders because the document structure will be familiar to them.

Make sure that your stakeholders sign off on the operating model *before* your new contact center moves into production, and make it a requirement for your decision to go live. Each team member needs to understand their role and the process for operating the contact center in the production environment.

Service introduction (SI)

SI activities implement the changes defined in the operating model. Think of the operating model definition as the design and build phases for the new model, and SI as the deploy phase of the operating model.

The SI team is often a dedicated team in your organization, and works independently from the project team. The project must pass the SI team's criteria and checklists before it receives approval to go live. For example, checklists include user acceptance testing (UAT) results and confirmation that a clashing event (such as a change freeze or another scheduled go-live event) isn't taking place on the same day that the project goes live, that users have the necessary training, and that operations teams are ready to proceed.

Do not leave SI activities to the end of the project. Engage the SI team early in the project and include them in your distribution list for design documentation. Early involvement ensures that the SI team can assist in the preparation to go live, such as helping to select the most suitable AWS support plan, providing impact statements for change requests (CRs), and supporting change approval board (CAB) discussions.

Training

Creating training materials and conducting well-attended training sessions are vital to successful migrations. Technology can work perfectly, but if users don't know how to answer calls and perform their daily tasks, the migration will be considered a failure.

Training activities might include direct user training, training the trainers, training for supervisors, training for support staff, and training for system administrators or product owners. Each organization is unique, so some options might be a better cultural fit than others.

We recommend a *train the trainer* approach that involves your organization's in-house training staff. Your staff knows your organization's culture, and the training format and techniques that best suit your users. Project team members can take on subject matter expert (SME) roles to provide technical materials (such as user manuals, administrator console manuals, and screen guides) that can be used as source material for the train the trainer sessions. If your organization doesn't have a training team, project SMEs should train supervisors and lead support staff, who can then train the users of the contact center.

We also recommend that system administrators and product owners take formal, instructor-led product training courses to gain a deeper understanding of the AWS environment and Amazon Connect console, so they can use product features and troubleshoot effectively.

Technical foundation workstream

This workstream involves decisions that require significant rework if changed, so the workstream emphasizes careful design, wide consultation, and up-front investment in DevOps processes and testing.

The technical foundation workstream consists of five phases: discovery and roadmap, design, build, test, deploy, and post go-live support.
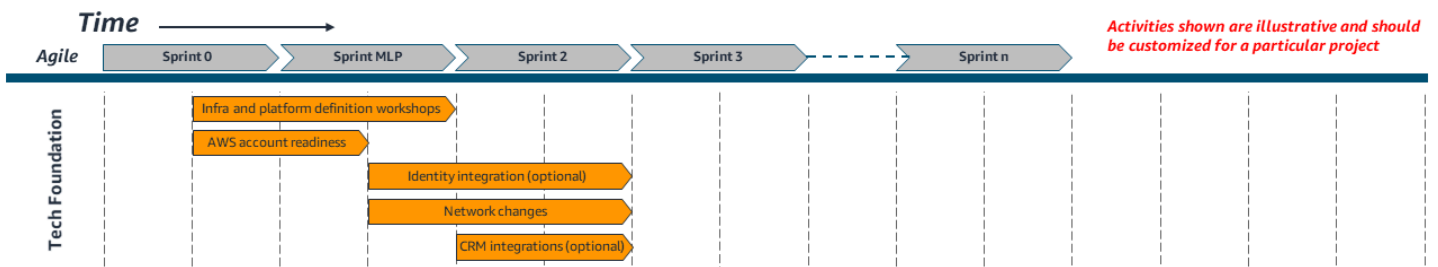
Discovery and roadmap

In this phase, you gather information and schedule workshops for the following:

- As-is mapping – Examine systems and capabilities, gather data, and meet with SMEs to understand the current state of the contact center.

- To-be design and gap assessment – Determine the ideal experience for all contact center agents and customers to determine project scope.

- Gap closure plan – Outline a roadmap for building and deploying the future state of the contact center.

Workshop attendees:

- Project managers

- Business, solutions, technical, and security architects

- Infrastructure platform owners

Design

In this phase, you produce design documents. You might have your own conventions or processes for creating design artifacts. We recommend including at least three sections in the design document: Amazon Connect configuration, networking, and security. Each section will likely have different, specialized stakeholder groups to ensure effective reviews and sign-offs, so it might be more practical to create separate documents for these three areas. Stakeholders should include architects, the security and compliance team, and platform owners.

Build

In this phase, you follow infrastructure as code (IaC) principles by using DevOps tools to standardize and manage stable releases. Avoid adopting a manual build process, even if it helps you get started more quickly, because this can increase the risks to stability and the number of bugs as the build becomes more complex and is promoted to the test and production environments. If you don't have your own DevOps tools, we recommend that you use AWS tools such as AWS CodePipeline and AWS CodeBuild, which can be switched on quickly. Factor the effort for setting up these tools into the scope of the project; they will be beneficial in the long term and enable you to follow DevOps principles. We recommend that you build in at least three separate AWS accounts for development, testing, and production. DevOps tools and automation can help you move code through these environments.

Test

The test phase consists of three sequential subphases:

1. Unit testing – Testing individual infrastructure components to ensure that they are correct and within design specifications. Performed by: developers

2. Integration testing – Testing items that form integration boundaries, such as Microsoft Active Directory (AD) identity management services. Performed by: developers

3. Product testing – End-to-end testing of functional journeys throughout the infrastructure; for example, testing that each agent event is logged in the security monitoring tool, the call is taken, and the call recording is in the correct Amazon Simple Storage Service (Amazon S3) bucket. Performed by: functional test team

Deploy

The infrastructure must be ready to handle live traffic when user journeys are scheduled to go live. The focus in the deployment phase is to ensure that AWS service quotas meet expected call volumes and the number of concurrent agents, number porting or toll-free number service (TFNS) repointing is complete, and the health of backend systems is being monitored as live traffic volumes ramp up. The security and compliance team should also confirm that the platform is ready for live traffic from their perspective.

Post go-live support (PGLS)

The project team remains engaged with the business as usual (BAU) support teams and end-users during the first few weeks after the new contact center goes live. The project team can help users get started on the new system, get involved in troubleshooting issues alongside the BAU support team, and improve support documentation based on feedback.

User journeys workstream

The user journeys workstream also involves decisions that can be changed or reversed with little effort or impact. The emphasis is on starting with a basic build of the user journey and iterating frequently and rapidly to get to the final product. It is rare for the final user journey to look exactly like the first one proposed, so an agile and iterative approach makes sense for this workstream.

The user journeys workstream consists of five phases: discovery, design, build, test, deploy, and post go-live support.

Discovery

In this phase, you gather existing user journey flows and designs, and pass them to the contact flow build team. If these do not exist or you want to design a new user journey, gather stakeholders in a workshop and collaboratively develop a user journey framework in a visual capture tool such as the following:

- Visual canvas tool – Use a tool such as Microsoft PowerPoint, Microsoft Visio, or draw.io. Screen-share the canvas to all stakeholders in a workshop. Add blocks and decision points to build an end-to-end user journey, and add placeholders for steps that should be confirmed later (for example, the exact wording or import of a queuing message audio file). Add the name of the owner who should confirm the placeholder.

- Contact flow designer – Instead of using a drawing tool like draw.io or Visio, consider using the contact flow designer that's included in Amazon Connect to develop and document the user journey over screen-share. Use prompt block placeholders for steps that should be confirmed later (for example, the exact wording or import of queuing message audio file). Use a simple text-to-speech (TTS) prompt block to record the owner who is confirming the step (for example, "Queue A message .wav file to be provided by John Smith"). This enables you to perform end-to-end testing of the user journey and routing logic in parallel.

Workshop attendees:

- Project managers
- Business and solutions architects
- Business analysts
- Service line owner and operator

Design

Design documentation is optional. It depends on the size and complexity of the contact flow. If you use the contact flow designer, which has an intuitive, easy-to-follow flowchart interface, the journey is self-documented and represents the actual build of the contact flows. This ensures a single source of truth during the rapid and agile development of the user journey. Otherwise, standalone design documents for contact flows must adhere to change control to avoid diverging from the actual build over time.

Build

Amazon Connect configuration is available by using [AWS CloudFormation templates and APIs](#) in infrastructure as code (IaC) tools. Use DevOps tools to build and manage Amazon Connect components such as security profiles and contact flows. If you design flows by using the contact flow designer, you can include the flows in your IaC DevOps tools and manually export them as JSON files.

> ⓘ **Note**
>
> You can also start building contact flows in a development environment while other AWS accounts are being created, and export the flows into the test and production environments when their Amazon Connect instances are ready.

Test

The test phase consists of two sequential subphases:

- Functional testing – Performed iteratively over agile sprints as contact flows are created in Amazon Connect. Performed by: functional test team
- User acceptance testing (UAT) – Performed only after contact flows have passed functional tests. Performed by: client business users (a dedicated team or users from the service line business unit)

Deploy

In this phase, agent and user credentials are uploaded into the Amazon Connect production instance so that users can log in. You should upload contact flows only after they successfully pass UAT testing in the previous phase. Claim a temporary phone number in the Amazon Connect dashboard and assign it to the contact flows. These phone numbers will be visible only to the project team, who will use them to place test calls. The project team often runs a selection of UAT scripts during this process. This approach provides preparatory (*pipe-clean*) testing of the user journey before the system goes live and real agents can access the workflow. At the scheduled go-live time, this temporary number is replaced by the publicly routable number used by customers—this is the point where you cut over to the new system. If necessary, you can roll back changes by swapping the number back to the legacy service line.

Post go-live support (PGLS)

The project team remains engaged with the service line stakeholders, the business as usual (BAU) support teams, and end users during the first few weeks after the new contact center goes live. The project team can help users get started on the new system, get involved in troubleshooting live issues alongside the BAU support team, and improve contact flows based on customer and agent feedback.

Running a pilot

Completing an end-to-end migration project for a small business area allows for fast deployment without the risk of large-scale business disruption. This experience builds confidence in the value proposition (capability, operations, and cost) for a relatively small outlay and can be used to justify a larger release of funds and resources for a full-scale project.

Pilots gather lessons for a full-scale deployment based on how end-users react to the new platform. They help stakeholders answer important questions with real-life data such as the following:

- Is the training we're providing suitable and sufficient?

- Do new processes work properly when end-users take real calls?

- Are users distracted by other applications on their device?

- Does an architecture or pattern work as expected in the live environment?

Best practices

- Ideally, pilots should become part of the initial minimum lovable product (MLP) delivery in an early sprint.

- Participants in a pilot should include technical users, business users, and end-users.

- Interview stakeholders to get anecdotal feedback on how they use the system, and capture data on average handle time, abandonment rate, and so on, to compare the new system with previous platforms.

- Make sure that tweaks and amendments that are identified during the pilot are tracked to completion.

- Define your success criteria and next steps before the pilot begins. Success criteria should be data-driven to enable conclusive scoring to reach a success/fail decision. If stakeholders sign off on the pilot and the delivery plan for any amendments, the predefined next step (for example, to start a full-scale deployment) is initiated.

- Be positive when your pilot reveals areas that have to be changed or even redesigned. This is a valuable outcome of the pilot and builds the foundation for a successful go-live deployment. Do not aim for a pilot with zero recommendations—this outcome would raise concerns on the validity of the pilot.

Selecting a pilot group

The business area you select to pilot the solution would ideally demonstrate all capabilities in scope of the minimum lovable product (MLP) to meet business outcomes. The successful delivery of the MLP becomes the starting point for building out complexity and adding service capabilities. The MLP pilot group should:

- Represent a non-critical business area (for example, an internal help desk, or a notification of change of circumstances).
- Handle a low volume of calls, so users have the time to learn the new platform and to record their feedback and observations.
- Be trusted by the project team and stakeholders, to ensure that feedback is fair, accurate, and objective. This helps instill confidence in the outcome of the pilot and helps create a collaborative development environment.
- Perform the majority of in-scope platform functions. There is little value or relevance in a pilot that uses only ten percent of the functions that are in scope of the full-scale deployment.
- Perform a function that might have been excluded from, or not fully integrated in, the old platform because of technical limitations (such as remote work) or licensing. By starting with a group that has no reports or recordings in the old system, you might be able to avoid building legacy integrations or migrating legacy data. However, you should make sure that the pilot continues to represent the full-scale deployment.

In reality, you might have to compromise on some of these factors, depending on the ability and willingness of teams in your organization to take part in a pilot.

Best practices for migrations

Migrating to Amazon Connect is likely to change your contact center's technical architecture and your staff's daily processes. To minimize disruptions, follow the best practices in this section when you design and build your new contact center.

- Technical considerations
- Operational considerations

Technical considerations

For more information about the following technical best practices and additional recommendations, see Best practices for Amazon Connect in the *Amazon Connect Administrator Guide*.

Voice traffic path – Will audio streams travel across your corporate internet link, or should you use an AWS Direct Connect connection as a dedicated link? AWS Direct Connect avoids latency-sensitive voice traffic competing with general traffic across data center internet pipes such as web browsing and email.

Setting up your network – A healthy end-to-end network connection is essential for a consistent and stable user experience. You should consider every component, from the agent's device, through their local network connection and virtual private network (VPN), if applicable, to Amazon Connect. A network connection is only as healthy as its weakest link. To optimize your network for Amazon Connect, review Set up your network in the *Amazon Connect Administrator Guide.*

Remote agents – Do your agents use a VPN when they work from home? If so, consider enabling VPN split tunneling for voice traffic. This routes delay-sensitive voice traffic across the local internet instead of sending it back to the data center and routing it out to Amazon Connect over the internet. If you don't use split tunneling, latency is needlessly increased (resulting in delayed audio or sluggish soft-phone actions), additional traffic load is placed on the VPN concentrator device, and your data center internet ingress and egress charges go up.

Data migration – For data such as call recordings and reporting statistics, consider two approaches:

- Migrate the data to the new platform. This requires planning and feasibility assessment (for example, to check audio format compatibility), but means that you can access your legacy data from a single portal on the new platform.

- Archive your data in place and decommission it when the minimum retention period expires. This might be more cost-effective, especially if data is stored on a purchased platform and accessed infrequently so that having two portals to browse old and new data is a practical option.

Number porting

- Consider whether a non-geographic number (NGN) or toll-free number service (TFNS) provider is required. Porting toll-free, local rate, or direct-dial-in (DDI) numbers to Amazon Connect allows for centralized management and billing of the end-to-end call. Consider the current charging model for your NGN/TFNS service and compare it with Amazon Connect charges. Be mindful of charges for calls that are made outside operating hours. Some NGN/TFNS providers do not charge for these calls if they handle the out-of-hours check and messaging. NGN/TFNS contracts and terms vary, so collect the information carefully to perform an accurate comparison.

- Timelines for number porting can take several weeks, so file the porting request through a ticket as early as possible. Use the ticket to finalize a cutover date and time. If there are challenges with the timeline, temporarily set a number forwarding transfer from your existing telephony queue to the new Amazon Connect phone number, as detailed in cutover option below.

Cutover approaches for porting numbers

You can use NGN backend repointing or number porting to port phone numbers.

NGN backend repointing – Perform a backend repoint of the frontend NGN number to the inbound number (DDI) hosted on Amazon Connect, as shown in the following diagram. This does not require any public-facing number changes and is typically managed as a service request ticket to the NGN carrier provider. Repointing can be scheduled for a specific date and time.

Number porting – This process consists of two stages:

- Number forwarding – This optional step, illustrated in the following diagram, directs traffic from the old platform to the new platform without changing the public-facing number. You can complete this step before your scheduled number porting date. This expedites the migration of agents onto the new platform in parallel with the number porting process. It also allows for rapid rollback (which depends on a relatively simple change to call forwarding rules) without any dependencies on a carrier. However, we recommend not leaving number forwarding in place for a long period of time, because it increases call charges (you pay for inbound traffic on DDI-1, outbound forwarding, and inbound traffic on the new DDI-2) and consumes infrastructure capacity (each inbound call also consumes an outbound circuit for the forwarding path).

- Completion of number porting – On an agreed date and time, the carrier for DDI-1 ports the number to AWS, so it becomes available for Amazon Connect to use, as illustrated in the following diagram. You can then assign the number to user journeys or functions, and manage it as if it were a natively sourced DDI in AWS. This simplifies billing and provides flexibility, because you can manage telephone numbers in the Amazon Connect console instead of relying on a third-party carrier to process service requests.

Transferring calls between other platforms and Amazon Connect – Organizations often migrate agents to Amazon Connect in groups based on line of business, job type, or other criteria. During a period of time, agent groups on other platforms are progressively migrated to Amazon Connect. Depending on the number and size of the groups, the migration phase might take several months, and teams that are spread across different platforms might have to transfer calls to each other during this period.

To transfer calls between platforms, use PSTN DDI numbers. Assign these DDIs only for cross-platform transfer usage, so you can measure and report on transfers independently, and prioritize calls differently if required.

Consider whether call attached data has to be exchanged between platforms during transfers. For example, if a caller has passed security checks on one platform, their security status should be exchanged during the call transfer to prevent them from having to go through security again with an agent on Amazon Connect. There are two approaches to consider:

- **Transfers without call attached data** – Structure migration group phasing to reduce the operational need for transfers where call attached data would be necessary. For example, migrate teams that frequently transfer calls to each other together, after a caller has exchanged a significant amount of data, which would otherwise need to be recaptured. If a caller interacts only minimally with IVRs or agents before being transferred across platforms, it might be unnecessary to exchange call attached data. You should also consider expediting migration timelines to minimize the period when cross-platform transfers would be performed. This means accepting a temporary inconvenience in exchange for not having to build technical debt and manage a cross-platform data exchange solution that will no longer be required after migrations are complete.

- **Transfers with call attached data** – This approach is relevant for teams that will be spread across platforms for a significant period of time and need to have call attached data exchanged during transfers to maintain operational performance. Use a technique called *rolling dialed number identification service (DNIS)*. For an example of how you can get started with rolling DNIS, see the GitHub repository Transfers from Legacy Platform into Amazon Connect.

Separate AWS accounts – Set up different AWS accounts for your Amazon Connect development, testing, and production instances. This approach separates those activities and limits the impact of changes to a single account. It also provides billing boundaries so that the appropriate business unit can pay for development, testing, and production work.

You can create new accounts with specific policies, rules, and principles, based on predefined templates. This means that any build or configuration in that account will have to comply with the specifications defined by the organization. You can use AWS Organizations to centrally manage and govern accounts.

Logging and alerting – Enable Amazon CloudWatch Logs to track usage thresholds and errors in contact flows. You can view usage and errors by using CloudWatch dashboards. You can also proactively send alerts through email or SMS text messages. By gaining visibility into low-level system behavior, you can identify and resolve issues quickly before they become bigger problems. An example of a proactive alerting solution for Amazon Connect is described in the blog post Monitor and trigger alerts using Amazon CloudWatch for Amazon Connect.

Single sign-on (SSO) – Use SSO to enable users to log in to Amazon Connect by using their corporate credentials (for example, through Active Directory) instead of requiring a separate username and password. This provides the optimum user experience because it doesn't require an additional login step or another set of credentials. It also avoids the need to centrally manage separate login credentials for password resets and other operations. Amazon Connect supports a number of identity management integration patterns. For more information, see Plan your identity management in Amazon Connect in the *Amazon Connect Administrator Guide*.

Workstation devices – Verify that end-user (for example, agent and supervisor) machines meet the minimum CPU and memory requirements noted in the Agent headset and workstation requirements for the CCP section of the *Amazon Connect Administrator Guide*. If you're planning to use these workstations for tasks outside contact center work, they should meet higher requirements. Use the Amazon Connect Endpoint Test Utility to check device and network compatibility. We recommend that you run this utility on a variety of agent workstations at

different locations, including agents who are working from home or from distinct network-island locations, to ensure compatibility across your organization.

Virtual desktop infrastructure (VDI) environments – Consider [network](#)and [deployment](#)optimizations for virtual desktop users.

Headsets – Use wired USB-powered headsets to ensure a consistent audio experience. Discourage the use of bluetooth or wireless headsets, which can add latency and reduce audio quality.

Wired network connections – Devices should use wired (Ethernet) connections to ensure a stable, high-quality audio experience. Verify that devices have wired ports. If a dongle is required, it must be budgeted and procured before migration.

Microphone and speaker settings – If your organization uses multi-purpose devices, confirm that shared use of microphones and speakers is allowed (turn off exclusive mode). For guidance, see [One-way audio from customers?](#) in the *Amazon Connect Administrator Guide*. This guidance applies to both speakers and microphones.

Dedicated devices (ideal) – If possible, users should be given devices for exclusive contact center use. You can then optimize these devices for the contact center experience, and use different devices for other duties.

Legacy habits – Watch out for legacy user behavior that might affect new processes. For example:

- Do agent devices predominantly connect over Wi-Fi today? If so, requiring wired connections will be a cultural shift for agents and might lead to poor compliance and call experience. An end-user education campaign might be required to drive this culture shift.

- Do agents use other collaboration applications (such as Microsoft Teams or Zoom) on their devices? This can lead to conflicting demands for speaker and microphone devices on the device, such as when Amazon Connect tries to deliver an incoming call while the agent is on another call. It can also lead to agents missing customer calls because they are busy making internal calls. We recommend removing other collaboration applications, where practical, to avoid call clashes.

Operational considerations

The best practices in this section focus on smoothing operations and keeping end-users happy with the new contact center platform and processes so they can provide constructive feedback. If end-users feel ignored or under-valued during the project, they will be reluctant to move to the new

platform. If end-users are unhappy, the migration will be considered a failure regardless of how well the technology performs.

Shifting to soft phones – Will this be the first time agents use a soft phone, which provides the phone interface on the screen, because they currently control calls through a physical desk phone? If so, it might be difficult for agents to transition from pressing buttons on a desk phone to using a soft phone keypad on a PC.

- Make sure that adjustment time is included in the training schedule. Expect a learning curve after the new contact center goes live.

- Accessibility might be a concern for agents who are used to desk phones, which are tactile devices. Consult agents who have accessibility concerns and include their feedback in design specifications for the soft phone color scheme and keypad button sizes.

Desk phone alternative – Agents can have calls delivered to a desk phone, as explained in the Amazon Connect setup instructions, as an alternative to a soft phone. This alternative handset must have a publicly reachable phone number, which is then configured in the agent profile. For example, this can be useful when a remote internet connection cannot support high-quality audio on the soft phone audio. In this case, the audio is sent across the traditional (PSTN) phone network.

Device inventory – Ensure that end-users have the right equipment on the day the new contact center goes live:

- Desk phones are no longer needed, so they can be decommissioned to free up desk space.

- Devices (such as laptops) might need Ethernet dongles to support hardwired Ethernet connections. Issue these to users before the go-live date to avoid last-minute demands that will affect your local IT parts team.

- Devices might have to provide faster CPU and more memory to run soft phone and business applications in parallel. Perform real-life testing (during UAT) with end-users by using the soft phone *alongside* their usual applications, to see if devices remain performant.

Support model (raising support tickets, levels 1-3 technical support desk ownership) – Work with your AWS account team, such as your technical account manager (TAM), to verify that you're on the most suitable AWS support plan. Make sure that everyone knows their role in the support model—from receiving incident reports from end-users to raising incident bridges for business-critical issues. Simulate an issue by raising a test incident to the level 1 support desk and tracking

it through support model processes. This will help you find gaps in the support model, so you can avoid issues after you go live.

Back office – Consider how tasks will flow between front-office agents and back-office teams. For example, the process for transferring calls and escalating customer cases might change. Include task workflows and routing in your test scripts.

Billing – AWS billing costs will increase and legacy platform costs will decrease immediately after the new contact center goes live. Contact center charges will be subsumed into AWS billing after migration. Notify your finance and accounting teams of this change so that costs from AWS accounts that host Amazon Connect instances can be mapped to the appropriate business unit. This is likely the same business unit that's responsible for legacy platform charges.

Access permissions – You can provide granular permissions to your contact center users by creating security profiles in Amazon Connect. This feature enables you to create advanced user access models based on the principle of least privilege to perform a role. On legacy platforms, permissions are typically granted too broadly. In contrast, in Amazon Connect, you can give users access to very specific resources and activities. For example, you can grant employees permission to edit users but not create or delete them, or to view user journey contact flows but not change them. Granular permissions are a powerful way to improve user engagement and optimize how responsibilities are distributed across roles and (such as agents, operators, supervisors, and developers) and teams. In addition to using security profiles, you can use Amazon Connect with AWS Identity and Access Management (IAM) features and polices. For more information, see How Amazon Connect works with IAM in the *Amazon Connect Administrator Guide*.

Service quotas – Service quotas are default settings that protect you from unexpected load and consumption charges. For example, service quotas can limit you to 10 concurrent calls or 5 phone numbers per instance. We recommend that you view your service quotas and request increases to support your expected usage. For more information, see Amazon Connect service quotas in the *Amazon Connect Administrator Guide*.

Agility through DevOps – Use a DevOps deployment pipeline to accelerate your release schedules and deliver new features more frequently. Business owners might have to reset expectations on how quickly they can release software, because the technology is more agile. Using deployment pipelines enables you to release smaller code bundles more frequently, so your releases are less risky and reach your customers faster.

Migration checklists

Use the following checklists to ensure that you complete important migration activities in the correct order.

Before you go live

1. Verify that the release passes user acceptance testing (UAT) and that any remaining issues have been accepted by stakeholders.

2. Plan the phone number cutover:

   - If you're using toll-free number service (TFNS): Verify that the service is ready to repoint to the Amazon Connect queue phone number. This might be a self-service task or might require a ticket with the provider, so consider the lead time to complete this task.

   - If you're porting the number to AWS; File a number porting request ticket well before the target go-live date. (See *Number porting* in the [Best practices for migrations](#) section earlier in this guide.)

3. Verify that end-users have been trained and know how to use the new platform.

4. Verify that the operations team has signed off on the new platform and has incorporated it into their support model. For example, the business as usual (BAU) team should be ready to manage any support tickets that are opened on the new platform.

5. Verify that the code base has been deployed to the production environment.

   > **ⓘ Note**
   >
   > This activity might require its own change request (CR), which would be submitted before and separately from the go-live CR for cutover.

6. Verify that the service lines that are in scope have successfully run UAT scripts by using a temporary phone number.

7. Submit a change request (CR) for go-live cutover and gain approval from the relevant change approval board (CAB). The evidence from this checklist is provided as input into the CAB discussion. The outcome of the CAB discussion is an approval to perform a cutover on a specific date and time.

On the day you go live

1. Ensure that agents are logged in to Amazon Connect and are available to receive and make calls and participate in chats. Supervisors and operators can check agent activity by using real-time reports on the Amazon Connect dashboard.

2. Ensure that the post go-live support (PGLS) team is present and ready.

3. (Optional) Confirm that staff who can assist agents and help troubleshoot problems is in position (on-site or at remote help desk).

4. Ensure that BAU support teams are aware of the cutover time and are ready to handle any support tickets.

> ⓘ **Note**
>
> The PGLS team works alongside the BAU support teams.

5. Open a conference bridge for stakeholders to receive status updates. This bridge also serves as a forum to discuss any issues that might occur. Keep the bridge open until go-live (or rollback) activities have been completed successfully.

6. Initiate the cutover (for example, the TFNS repoint) at the approved time.

7. Review real-time metrics on the Amazon Connect dashboard to verify the following:

   - Calls are being answered.

   - Abandonment rates and average handle times (AHT) are as expected.

   - Queue depth remains sensible.

Post-migration optimizations

Your work to develop and improve the user experience doesn't end on the day you go live. Amazon Connect and AWS have tools that provide detailed business insights, from granular reporting to fraud detection and voice biometrics driven by artificial intelligence (AI). This information helps you add new and innovative capabilities, and transform the customer and agent experience in your contact center.

You can use agile delivery methods to deliver new capabilities as sprint iterations after you go live. You can prioritize new capabilities and optimizations, and add them to a sprint backlog.

Examples of innovative capabilities that help deliver significant changes in operations and user experiences include the following:

- Amazon QuickSight dashboards provide easy-to-use metrics and graphical reports, and enable supervisors to monitor agent utilization to ensure balanced staffing across teams.

- Proactive alerting through email and SMS when defined operational thresholds are breached helps you identify problems before an issue or outage occurs. For example, if queue depth or average handle time (AHT) values rise above a defined limit, proactive alerting enables supervisors to rapidly intervene.

- Contact Lens for Amazon Connect performs sentiment analysis by using AI and speech recognition to transcribe a call. It can generate alerts on profanity or negative sentiments, and enable supervisors and agents to escalate these problems.

- Amazon Connect Voice ID provides voice biometrics based on a few seconds of speech audio. This voice print can be gathered during normal conversation with a bot or agent and eliminates having to remember pass phrases and secret answers. This feature greatly improves customer experience and reduces agent handle time.

- Amazon high-volume outbound dialer provides a a way to reach millions of customers to communicate news, reminders, and delivery notifications without requiring any third-party tools. This feature automates dialing and includes voice mail detection to connect agents to real customers with minimal effort, without having to look up customer records manually.

- An array of AWS-powered data analytics, AI, and machine learning (ML) tools are available, including Amazon Athena, Amazon Comprehend, and Amazon SageMaker AI. Apply models to look for trends in interactions that could lead to business insights, such as:
  - Fraud detection

- Frequent utterances, to identify what people are calling about, possibly leading to proactive messaging campaigns or contact center team changes

- High-touch customers who call more frequently than others, possibly allowing targeted outreach by an agent to preempt them from calling in

A successful migration is only the start of the journey to reimagine and transform your contact center. AWS services provide innovative experiences that you can add to your contact center to generate unique customer and agent experiences.

Next steps

If you're planning to migrate your contact center to the cloud, you might be concerned about how the migration will affect your customer portal and your brand. If you have the right vision, a robust delivery plan, and continued innovation after you go live, the migration can be a success from multiple angles: technical, operational, and financial.

Include some form of transformation in the initial phase of your migration plan to improve the customer experience. Establish mechanisms to work back from the customer and to listen to the voice of the customer to fuel this innovation. Use real data and end-user insights as much as possible. Ultimately these innovations will reduce customers' efforts to resolve issues and will increase customer retention and loyalty.

This strategy is a starting point for planning your migration journey. Reach out to your AWS account manager or fill out the [AWS Professional Services form](#) for more information or if you need help in any of these areas:

- Resource constraints
- Help developing AWS competencies and skills
- Help working with agile methodologies
- Time constraints, need for acceleration

Resources

Books

- Dixon, Matthew, Nick Toman, and Rick DeLisi. 2013. The Effortless Experience: Conquering the New Battleground for Customer Loyalty.

Case studies

- The Amazon Connect Customers website has a list of case studies categorized across industries.

Partners

- Amazon Connect Delivery Partners are AWS Partners who help companies build cloud contact centers with Amazon Connect. These AWS Partners can help you improve customer experiences and outcomes through Amazon Connect.

Official blog

- AWS Contact Center Blog hosts articles written for business and technical users. Use these to discover market insights, new ideas, and ways to optimize your contact center.

AWS Online Tech Talks

- Migration Best Practices and Resources: Moving Your Contact Center to Amazon Connect

Useful links

- AWS Migration Acceleration Program (MAP)
- AWS Cloud Adoption Framework (AWS CAF)
- AWS Professional Services (contact AWS Sales from this page)
- AWS Prescriptive Guidance
- Amazon Connect Administrator Guide
- Amazon Connect resources

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
Cross-platform call transfers	Expanded the informati on about transferring calls between other platforms and Amazon Connect.	December 6, 2024
New best practice	Added information about DNIS to the Technical considerations section.	November 11, 2024
Initial publication	—	August 24, 2022

AWS Prescriptive Guidance glossary

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

Numbers

7 Rs

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

- Refactor/re-architect – Move an application and modify its architecture by taking full advantage of cloud-native features to improve agility, performance, and scalability. This typically involves porting the operating system and database. Example: Migrate your on-premises Oracle database to the Amazon Aurora PostgreSQL-Compatible Edition.

- Replatform (lift and reshape) – Move an application to the cloud, and introduce some level of optimization to take advantage of cloud capabilities. Example: Migrate your on-premises Oracle database to Amazon Relational Database Service (Amazon RDS) for Oracle in the AWS Cloud.

- Repurchase (drop and shop) – Switch to a different product, typically by moving from a traditional license to a SaaS model. Example: Migrate your customer relationship management (CRM) system to Salesforce.com.

- Rehost (lift and shift) – Move an application to the cloud without making any changes to take advantage of cloud capabilities. Example: Migrate your on-premises Oracle database to Oracle on an EC2 instance in the AWS Cloud.

- Relocate (hypervisor-level lift and shift) – Move infrastructure to the cloud without purchasing new hardware, rewriting applications, or modifying your existing operations. You migrate servers from an on-premises platform to a cloud service for the same platform. Example: Migrate a Microsoft Hyper-V application to AWS.

- Retain (revisit) – Keep applications in your source environment. These might include applications that require major refactoring, and you want to postpone that work until a later time, and legacy applications that you want to retain, because there's no business justification for migrating them.

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

A

ABAC

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

abstracted services

See [managed services](#).

ACID

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

active-active migration

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

active-passive migration

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

aggregate function

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

AI

See [artificial intelligence](#).

AIOps

See [artificial intelligence operations](#).

anonymization

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

anti-pattern

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

application control

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

application portfolio

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

artificial intelligence (AI)

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

artificial intelligence operations (AIOps)

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

asymmetric encryption

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

atomicity, consistency, isolation, durability (ACID)

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

attribute-based access control (ABAC)

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

authoritative data source

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

Availability Zone

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

AWS Cloud Adoption Framework (AWS CAF)

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

AWS Workload Qualification Framework (AWS WQF)

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

B

bad bot

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

BCP

See business continuity planning.

behavior graph

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

big-endian system

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

binary classification

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

bloom filter

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

blue/green deployment

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

bot

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

botnet

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

branch

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

break-glass access

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

brownfield strategy

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

buffer cache

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

business capability

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

business continuity planning (BCP)

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

C

CAF

See [AWS Cloud Adoption Framework](#).

canary deployment

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

CCoE

See [Cloud Center of Excellence](#).

CDC

See [change data capture](#).

change data capture (CDC)

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

chaos engineering

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

CI/CD

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

classification

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

client-side encryption

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

Cloud Center of Excellence (CCoE)

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

cloud computing

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

cloud operating model

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

cloud stages of adoption

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

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

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

- Migration – Migrating individual applications

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

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

CMDB

See configuration management database.

code repository

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

cold cache

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

cold data

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

computer vision (CV)

A field of AI that uses machine learning to analyze and extract information from visual formats such as digital images and videos. For example, AWS Panorama offers devices that add CV to on-premises camera networks, and 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.