



Unable to locate subtitle

Migration Launch Guide



Migration Launch Guide: ***Unable to locate subtitle***

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

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

Table of Contents

Migration Launch Guide	1
AWS Transform	3
Services and agents	3
AWS Transform assessments (migration assessments)	3
AWS Transform for VMware migrations	4
AWS Transform for Windows modernization	5
AWS Transform for mainframe modernization	6
AWS Transform custom (code modernization)	6
Latest innovations	7
Benefits	8
How to access AWS Transform	8
Customer success	9
Getting started	10
Migration Acceleration Program (MAP)	11
What is tagging for MAP?	11
General best practices for tagging resources on AWS	12
MAP tagging best practices	12
MAP 2.0 customer-facing dashboard	13
Monitor tagged or untagged services eligible for MAP credits	13
Delivering a migration program	15
Migration planning	15
AWS Transform	15
Staffing the migration	16
Cloud Center of Excellence (CCoE)	16
AWS partners	16
AWS Professional Services	17
AWS Countdown Premium	17
AWS Managed Services	17
Migration tooling	17
Case studies	18
Whitepapers and AWS prescriptive guidance	18
AWS Well-Architected Labs	19
Modernization pathways	20
Move to AI	20

Move to cloud native	21
Move to containers	22
Move to managed databases	23
Move to open source	24
Move to modern analytics	25
Move to modern DevOps	26
Generative AI insights for modernization pathways	28
How modernization enables Gen AI	28
Gen AI insights by pathway	29
Move to AI: Agentic AI and autonomous workflows	29
Move to cloud native: Intelligent microservices and event-driven AI	30
Move to containers: Scalable AI inference and model serving	31
Move to managed databases: RAG, semantic search, and data intelligence	31
Move to open source: Flexible model selection and fine-tuning	32
Move to modern analytics: Generative BI and predictive insights	32
Move to modern DevOps: AI-powered operations and developer productivity	33
Summary	35
Organizational readiness	37
Cloud Center of Excellence (CCoE)	16
AWS Cloud Adoption Framework (CAF)	38
Cloud Readiness Assessment (CRA)	38
Cloud Maturity Assessment (CMA)	38
Cloud Operations	39
Maturity of cloud operations	39
Additional resources	39
Cloud financial management and FinOps	39
Operational readiness	41
Managing your resources	41
Monitoring and observability	41
Security and compliance	42
Backup	43
Optimization and cost management	43
Education and enablement	44
Learning-needs analysis	44
AWS training and certification	44
AWS Skill Builder	44

AWS GameDay	45
AWS Workshops	45
Experience-Based Acceleration	45
Getting support	47
Resources	48
Migration Acceleration Program (MAP)	48
Delivering a migration program	48
Organizational readiness	50
Operational readiness	51
Education and enablement	52
Generative AI and modernization	52

Migration Launch Guide

This guide offers guidance on the resources that customers commonly use when [migrating to or modernizing on](#) AWS. The guide is not intended to be a prescriptive implementation guide for cloud migration or modernization, after all *'There is no compression algorithm for experience'* (Andy Jassy, Amazon CEO). The guide has been prepared by an experienced team of AWS Customer Solutions Managers that actively support customers in their AWS migration journey.

The Migration Launch Guide acknowledges that no two cloud journeys are the same. So it provides guidance on both technical and non-technical topics and considers the various degrees of cloud maturity an organization may have. If you are participating in the [Migration Acceleration Program \(MAP\)](#), you will find useful information regarding both the program and its onboarding requirements, to help you to fully leverage its advantages.

The section on [AWS Transform](#) introduces the agentic AI services and specialized agents available to accelerate your migration and modernization journey—from assessments and VMware migrations to Windows, mainframe, and custom code transformations.

The section on [Delivering a migration program](#) describes the phases of a migration journey. It gives guidance on how to plan your migration, considerations for staffing and allows you to get acquainted with some of the tools you can use to simplify, automate and guide you through your migration.

The section on [Modernization pathways](#) outlines structured approaches to modernizing application portfolios after migration. Most organizations choose a combination of pathways to meet specific business outcomes.

After implementing a modernization pathway, the [Generative AI insights for modernization pathways](#) section maps the insights and use cases that become available as a natural next step in the journey.

As you progress your [Organizational readiness](#) by using AWS, you can find guidance on how to determine where you are on your cloud journey. The tools that help identify potential gaps in your organizations' plan, skills and governance, and allow you to focus on what is important to you.

To get a head start on cloud operations, see [Operational readiness](#).

The section on [Education and enablement](#) presents the most useful mechanisms to enable your organization to fully leverage AWS tools and services.

If you need additional support or a deep-dive into any discussed topic in this Migration Launch Guide, see [Getting support](#).

Finally, we have collated all relevant links from this guide into the [Resources](#) section.

AWS Transform

AWS Transform is a collaborative enterprise IT transformation workbench powered by specialized AI agents, agentic workflows, and continuous learning that accelerates cloud migration, legacy application modernization, and tech debt reduction. It offers purpose-built transformations for mainframe, VMware, and Windows workloads, and automates code modernization at scale—from version and framework upgrades to fully custom, organization-specific transformations.

By automating complex tasks like assessments, code analysis, refactoring, dependency mapping, and transformation planning, AWS Transform dramatically reduces project timelines and costs. Its shared workspaces and natural language chat experience enable cross-functional teams to collaborate in real time, track progress, and manage transformations from start to finish.

In its first year, AWS Transform has migrated tens of thousands of VMs and processed over 4.5 billion lines of code, saving 1.69 million hours of manual effort across mainframe, Windows/.NET, infrastructure, and custom code modernization.

To get started, see [AWS Transform in the AWS Management Console](#). For the latest features, see [AWS Transform](#). For detailed guidance, see the [AWS Transform User Guide](#).

Services and agents

AWS Transform provides specialized AI agents for five primary use cases. Each agent is purpose-built on 20 years of AWS migration and modernization experience and is accessible from the web console, CLI, IDE (Kiro, VS Code), or through MCP (Model Context Protocol) integration.

AWS Transform assessments (migration assessments)

AWS Transform assessments help you evaluate the cost, feasibility, and business value of migrating on-premises infrastructure to AWS. Using AI-powered agents, assessments provide automated right-sizing recommendations, multi-scenario comparison, and interactive refinement through natural language chat.

Key capabilities include:

- **Cost estimation** – Get cost estimates for Amazon EC2, Amazon RDS for SQL Server, Amazon EBS, Amazon S3, Amazon FSx, and end user computing services.

- **Automated right-sizing** – Analyzes on-premises server specifications and recommends appropriately sized EC2 instances based on CPU, memory, and performance requirements.
- **SQL Server assessment** – Specialized analysis for SQL Server migrations to Amazon RDS or EC2, including BYOM/LI licensing options and Database Savings Plans (up to 20% savings).
- **Business value assessment** – Evaluates staff productivity gains, resilience improvements, and business agility beyond infrastructure cost savings.
- **Sustainability** – Estimates carbon footprint reduction from migrating to AWS.
- **Multi-scenario comparison** – Create multiple scenarios with different assumptions (pricing models, regions, instance families, licensing) and compare side by side with what-if analysis.
- **Chat-based interaction** – Refine assessments interactively through natural language. Supports rough estimation with limited data, adding inventory through chat, and modifying costs on the fly.
- **Deliverable generation** – Generate executive presentations (PPTX) and detailed reports (PDF) directly from assessment results.

AWS Transform accepts inventory data from multiple sources including the AWS Transform Discovery Tool, RVTools (VMware), CMDB exports, Migration Evaluator, partner discovery tools, and MPA format files.

Use assessments at the start of your migration journey to build the business case, inform planning, and gain executive approval.

To learn more, see [Migration assessments](#) in the AWS Transform User Guide.

AWS Transform for VMware migrations

AWS Transform for VMware migrations accelerates enterprise cloud migrations with purpose-built AI agents that automate the entire lifecycle—from discovery and wave planning through network conversion and cutover. By unifying teams, tools, and workflows in a single collaborative workbench, migrations that once took years now complete in months.

The service supports VMware, bare metal, Hyper-V, and database workloads. Key capabilities include:

- **Intelligent wave planning** – AI agents analyze application dependencies, business priorities, and technical constraints to group workloads into optimized migration waves. What once required weeks of manual effort now takes minutes.

- **Landing zone creation** – Automates the generation of Landing Zone Accelerator (LZA) configurations with secure, multi-account AWS environments, consistent governance, and security controls. Supports CloudFormation, CDK, Terraform, and LZA formats.
- **Automated network conversion** – Converts complex on-premises network configurations into AWS equivalents (VPCs, subnets, transit gateways) up to 80x faster than manual methods.
- **Rehost or replatform** – Manages the full rehosting lifecycle (agent deployment, test instances, cutover orchestration) or replatforms applications to containers on Amazon ECS or Amazon EKS in the same workflow.

The service supports first-party AWS agents, third-party partner agents, and custom Bring Your Own Agent workflows.

To learn more, see [AWS Transform for VMware migrations](#).

Watch the following video from re:Invent 2025 on how to complete a large-scale migration with agentic AI: [Reimagining large-scale migration planning with agentic AI](#).

AWS Transform for Windows modernization

AWS Transform accelerates full-stack Windows modernization by up to 5x across application, UI framework, database, and deployment layers. Using the specialized Windows modernization agent, teams can modernize .NET applications, SQL Server databases, and deployment processes—reducing operating costs by up to 70% by moving away from costly licenses.

Key capabilities include:

- **.NET framework modernization** – Upgrades legacy .NET Framework applications to .NET 8+ (cross-platform), enabling containerization and Linux deployment.
- **SQL Server migration** – Automates database modernization from SQL Server to open-source engines (PostgreSQL, MySQL) or AWS-managed databases.
- **UI framework upgrades** – Modernizes legacy UI frameworks (for example, Web Forms to Blazor or React).
- **Deployment modernization** – Converts traditional IIS deployments to containerized microservices on Amazon ECS or Amazon EKS.

In addition to the AWS Transform web console, a Visual Studio extension (VSIX) provides an integrated IDE experience for .NET modernization.

To learn more, see [AWS Transform for Windows](#).

AWS Transform for mainframe modernization

AWS Transform accelerates mainframe application modernization from years to months. The specialized mainframe agent streamlines the entire transformation process—from initial analysis and planning to code refactoring and application reimagining—reducing risk and costs.

Key capabilities include:

- **Codebase analysis** – Analyzes COBOL, PL/I, JCL, and other mainframe languages to map dependencies, detect missing artifacts, and assess complexity with visual dependency representations.
- **Automated refactoring** – Converts mainframe code to modern languages (Java, C#) while preserving business logic. Works with Claude Code and other agentic IDEs for the forward engineering phase.
- **Data migration** – Migrates mainframe data stores (VSAM, IMS DB, Db2) to cloud-native databases.
- **Testing and validation** – Automated equivalence testing to ensure functional parity between legacy and modernized applications.

To learn more, see [AWS Transform for mainframe](#).

AWS Transform custom (code modernization)

AWS Transform custom is the agentic AI service for eliminating technical debt at scale. It transforms any code pattern—from any language, framework, or runtime to any other—through both built-in (AWS managed) and custom transformation recipes.

Key capabilities include:

- **Built-in transformations** – Ready out-of-the-box for common scenarios:
 - Java, Python, Node.js version upgrades (Lambda and non-Lambda)
 - AWS SDK v1 to v2 upgrades
 - Java X86 to Graviton
 - Log4J to SLF4J
 - Angular to React / Angular version upgrades

- Spring Boot upgrades, JBoss to Spring Boot
- Comprehensive code analysis and tech debt reporting
- **Custom transformations** – Create your own for any-to-any use cases. Define transformation requirements in natural language, test on sample repositories, then apply at enterprise scale. Supports diverse scenarios from language translations (Progress 4GL to Java, C to Rust, COBOL to Java) to architectural changes (X86 to Graviton, microservices decomposition).
- **Continual learning** – The agent learns from each execution, becoming smarter when applied repeatedly to similar repositories.
- **Scale** – CLI-based interface that acts on one or multiple repositories. Integrates with CI/CD platforms for continuous modernization. Can run from laptop, EC2, ECS/EKS containers, or any pipeline.
- **IDE integration** – Available through Kiro Power (VS Code extension), as well as agent skills for Claude Code, Cursor, and Codex.
- **Modernization Analysis (MODA)** – Scans code for cloud-native maturity gaps and maps findings to AWS modernization pathways. Completes in 5–30 minutes per repository.
- **Agentic Readiness Analysis (ARA)** – Evaluates whether systems are ready to be safely called by AI agents, covering APIs, identity, state management, human-in-the-loop controls, and observability.
- **Campaign management** – Web experience for managing large-scale transformation campaigns, tracking progress across multiple repositories.

To learn more, see [AWS Transform custom](#).

Latest innovations

AWS Transform continues to expand capabilities. Recent launches include:

- **MCP integration** – Invoke transformation capabilities directly from orchestrators, IDEs, and coding environments using Model Context Protocol.
- **Kiro Power and agent skills** – Pre-built playbooks and individual tasks accessible from Kiro, Claude Code, Cursor, and Codex.
- **Web-to-IDE handoff** – Start modernization jobs in the web console, then hand off to developer IDEs for human-in-the-loop edits and refinement.

- **Agent builder toolkit** – Extend AWS Transform with custom agents using Kiro Power combined with the AWS Transform base agent and SDK.
- **Expanded regions** – Mumbai, Tokyo, Seoul, Sydney, Canada, and Europe.
- **CloudWatch metrics** – Monitor cost, success/failure rates, and agent minutes.

For the latest updates, see [AWS Transform – Latest Innovations](#).

Benefits

The following are the primary benefits of using AWS Transform:

- **Accelerate modernization** – Modernize Windows, mainframe, and VMware applications up to 5x faster with agentic AI-powered automation of analysis, planning, documentation, and transformation tasks.
- **Deliver at scale** – Transform hundreds of applications in parallel. AWS Transform automates high-effort, repeatable tasks so teams can deliver larger, more complex projects faster.
- **Reduce costs** – Remove legacy infrastructure, licensing, and modernization costs. Customers report up to 70% reduction in operating costs by moving away from costly licenses.
- **Specialized agents built on AWS expertise** – Achieve consistent, reliable outcomes using AI agents built on 20 years of AWS migration and modernization experience.
- **Unified collaboration** – Shared workspaces and natural language chat enable cross-functional teams (project managers, architects, developers, security engineers) to work together in real time.
- **Flexible tooling** – Use AWS Transform from the web console, CLI, IDE (Kiro, VS Code), or through MCP integration with your preferred orchestrators and coding environments.

How to access AWS Transform

The following table describes the interfaces you can use to access AWS Transform.

Interface	Description
Web console	AWS Transform console (https://console.aws.amazon.com/transform)—collaborative workbench for planning and tracking

Interface	Description
CLI	Command-line interface for automation, CI/CD integration, and batch operations
IDE (Kiro Power)	VS Code extension and Kiro IDE for developer-centric workflows
Visual Studio (VSIX)	Visual Studio extension for .NET modernization
Agent skills	Pre-built plugins for Claude Code, Cursor, and Codex
MCP	Model Context Protocol integration for orchestrators and custom toolchains

Customer success

The following are publicly available case studies of customers using AWS Transform.

Customer	Use case	Outcome	Reference
IDEMIA	.NET modernization	4x faster, 30% cost reduction	Case study
Thomson Reuters	Windows/.NET modernization	4x faster modernization	Case study
Bridgestone	Mainframe modernization	Completed in 7 months, 90% efficiency gains	Case study
CSL	VMware migration (5,000 servers, 29 data centers)	10x faster planning, 30% operational cost savings	Case study

Customer	Use case	Outcome	Reference
ADP	Compliance infrastructure	Reimagined with AWS Transform	Case study
Toyota	Mainframe modernization	Accelerated with generative AI	Case study
Vector	VMware migration	34% faster than traditional methods	Case study
Mercedes-Benz	SAP modernization	RISE with SAP on AWS	Case study

Getting started

To get started with AWS Transform, use the following resources:

- **Web experience** – Access [AWS Transform](#) in the AWS Management Console.
- **User guide** – See the [AWS Transform User Guide](#) for detailed instructions.
- **Blog** – Read [One year. 4.5 billion lines of code. 1.6 million hours saved](#) for an overview of impact and lessons learned.
- **Migration Launch Guide** – This guide provides context on the broader migration journey and how AWS Transform fits within it.
- **Partners** – Explore [AWS Transform Partners](#) for assistance with your transformation initiative.

For more information, contact your AWS account team.

Migration Acceleration Program (MAP)

The [AWS Migration Acceleration Program \(MAP\)](#) is a program to support customers as they migrate to the cloud. MAP is based on AWS's experience migrating thousands of global customers to the cloud.

Watch the **Introduction to MAP** video in any of the following languages: [\[English\]](#) [\[French\]](#) [\[Italian\]](#) [\[Spanish\]](#)

Getting started

If you are getting started with AWS Migration Acceleration Program 2.0, the following are the required steps to start generating MAP credits:

1. **MAP tagging setup:** Ensure the tag key that is required for MAP 2.0 is activated in **AWS Billing > Cost Allocation Tags**. For more information, see [AWS account setup](#) in the AWS Migration Acceleration Program 2.0 Tagging Guide and [Using AWS cost allocation tags](#) in the AWS Billing User Guide.
2. **Tag selection:** Choose the appropriate MAP 2.0 tag key/value combination. For more information, see [Tagging key combinations](#) in the AWS Migration Acceleration Program 2.0 Tagging Guide.
3. **Tagging migrated workloads:** See the [list of MAP 2.0 eligible services](#) and [the section called "MAP tagging best practices"](#).

What is tagging for MAP?

A tag is a key-value pair that you apply to an AWS resource to hold metadata about that resource. Tagging makes it possible for organizations to categorize and manage their cloud infrastructure easily at scale. You can also use tagging to create new resource constructs for visibility or control. For example, you can use tagging to group together resources that make up a workload. To obtain the full benefits of MAP, you must tag the migrated workloads. For more information, see [Tagging your AWS resources](#) and [Guidance for Tagging on AWS](#).

If you are a new MAP customer, see [AWS Migration Acceleration Program 2.0 Tagging Guide](#). For a list of MAP-credit-eligible services, see [What services are in scope for MAP 2.0?](#) in the AWS Migration Acceleration Program 2.0 Tagging Guide.

General best practices for tagging resources on AWS

AWS Documentation: [Tagging AWS Resources](#) .

Blog Post: [Maximizing resource tagging at scale and across teams for your migration to AWS](#).

The following training video explains how to tag resources in AWS.

MAP tagging best practices

- **Tagging with CloudFormation:**

- You can use CloudFormation Linter to enforce tags. You can create rules to enforce certain tag keys, and include linting steps in pipelines. For more information, see [AWS CloudFormation Linter](#) on GitHub.
- When you create or update an AWS CFN stack, set tags in the AWS Management Console or by using the AWS CLI. CloudFormation propagates these tags to all the AWS resources that you create in the stack. For more information, see [create-stack](#) and [update-stack](#) in the AWS CLI command reference.

- **Tagging with AWS Config:**

- To identify and tag MAP 2.0 eligible resources that were created after the MAP agreement was signed, use the [AWS Config Conformance Pack](#) . For more information, see [AWS Config](#).

- **Tagging with AWS Transform MGN (MGN):**

- With MGN, you can enable MAP tagging to automatically tag your launched instances. To enable MAP tagging, go to the MGN console and choose **Add MAP tag to Launched Instances**, then specify the MAP tag value that you want to use. MGN will automatically tag your launched instances with the key map-migrated and the value required for the MAP program. MGN only tags Amazon EC2 instances and Amazon EBS volumes.

- **AWS remediation options:**

- Use AWS Config to create rules that check resources for required tags and to continuously monitor your resources against those rules. Use it with [Amazon EventBridge](#) to trigger automated responses to missing tag values. For example, you can set it to send an email through Amazon SNS. For more information, see [Using Amazon EventBridge Scheduler with Amazon SNS](#) in the Amazon SNS Developer Guide.

- **Additional resources:**

- [Service control policies \(SCPs\)](#)

- [Enforcing tagging on resource creation](#)
- [Tag Policies](#)
- [Restricting allowed tag key/values](#)
- [Amazon Bedrock MAP Tagging using the AWS CLI](#)

MAP 2.0 customer-facing dashboard

The **MAP 2.0 Dashboard** is a self-serve mechanism for eligible MAP 2.0 customers to track the progress of their MAP migrations. This dashboard is available in the billing console for payer accounts that are part of MAP 2.0 agreements signed after November 2022. This dashboard is only available for customers that have received their first tranche of credits. For more information, and to learn how to access the dashboard, see the [MAP 2.0 Dashboard User Guide](#).

Monitor tagged or untagged services eligible for MAP credits

After you tag the target AWS resources, check to see how much these resources are used for each management account in your migration. You can use Cost Explorer to see the approximate amount of the target AWS resources. You can also use it to create a report on AWS services that are eligible for MAP credits but have not been tagged. To do so, select the filters below.

Report parameters	Filter	Setting
Time	Date range	Specify a date range
Time	Granularity	Specify a granularity
Group by	Dimension	Select Service, Linked account or Tag, as appropriate
Filters	Service	Exclude non-MAP-2.0-eligible services
Filters	Usage type	Exclude all "DataTransfer" values
Filters	Usage type group	Exclude all "Data Transfer" values

Report parameters	Filter	Setting
Filters	Tag	Choose to view either tagged spend or untagged spend. You can also leave it blank and then use the 'Dimension' to show the split of tagged versus untagged spend
Advanced options	Aggregate costs by	Amortized costs
Advanced options	Additional data settings	Untick all

For more information, contact your AWS account team.

Delivering a migration program

AWS recommends performing the migration process in three phases: Assess, Mobilize, and Migrate and Modernize. To learn about these phases and about tools and services that are available to you as you go through your migration, see [How to migrate](#). For more detailed guidance and to learn about migration strategies, guides, and patterns, see [AWS Prescriptive Guidance](#).

Migration planning

Long-running cloud migration programs require the coordination of several workstreams such as program governance, landing zone (an operative target environment with security controls), migration, and application portfolio. An application portfolio assessment (and subsequent migration planning) represents a foundational initial activity for any migration program. To learn about this process, see [Application portfolio assessment guide for AWS Cloud migration](#).

Explore the 3 phases of migration defined in the AWS migration process:

1. **Assess:** At the start of your migration journey, you assess your organization's current readiness for operating in the cloud. Most importantly, you identify the desired business outcomes and develop the business case for migration. To help you with this goal, AWS provides [AWS Transform assessments](#), an Agentic AI migration assessment to simplify and optimize your cloud migration. Discover cost optimization and performance opportunities while getting detailed financial modeling to help you confidently plan your migration and maximize potential savings.
2. **Mobilize:** As part of the mobilize phase, you create a migration plan and refine your business case. You address gaps that were identified in the assess phase to better prepare your organization with a focus on building your baseline environment (the "landing zone"), driving operational readiness, and developing cloud skills. Watch the following video about the Mobilize phase: [Learn how to mobilize for accelerated cloud migration - AWS Virtual Workshop](#).
3. **Migrate & Modernize:** The goal of the migrate-and-modernize phase is to design, migrate, and validate each application. To learn more, watch the following video from 2024: [MAP Overview | Amazon Web Services](#).

AWS Transform

[AWS Transform](#) for migrations accelerates enterprise cloud migrations with purpose-built AI agents that automate the entire lifecycle—from discovery and wave planning through network conversion

and cutover. By unifying teams, tools, and workflows in a single collaborative workbench, migrations that once took years now complete in months. It offers purpose-built transformations for mainframe, VMware, and Windows workloads, and automates code modernization at scale, from version and framework upgrades to fully custom, organization-specific transformations. By automating complex tasks like assessments, code analysis, refactoring, dependency mapping, and transformation planning, AWS Transform dramatically reduces project timelines and costs.

Staffing the migration

Migration programs can be delivered through in-house staff, external resources, or a combination of both. AWS cloud expertise is essential to ensure a successful migration, so it is critical to invest in this early on.

To get started, see [Education and Enablement](#).

Cloud Center of Excellence (CCoE)

Your CCoE is the starting point for preparing your organization for your migration, it's as relevant to an organization working alone or with a partner. You can choose to accelerate the creation and maturity of your CCoE with the help of a partner or through AWS Professional Services.

To learn about the importance of establishing a CCoE, watch the following video from 2024 AWS Tech Talks - [Build a cloud center of excellence | AWS Events](#).

To learn about the role of the CCoE in your organizational readiness, see the [Organizational readiness](#) topic in this launch guide.

AWS partners

The [AWS Competency Partner Program](#) is designed to identify, validate, and promote AWS Partners with demonstrated AWS technical expertise and proven customer success. The AWS Competency Partner Program contains an extensive list of partners and offerings with a proven track record of delivery in specific industries, use cases, and workloads. The following are some competencies that can be relevant to your migration:

- [AWS Managed Service Provider Partners](#)
- [AWS Migration Competency Partners](#)
- [AWS Microsoft Workloads Competency Partners](#)

- [AWS SAP Competency Partners](#)
- [AWS Oracle Competency Partner](#)
- [AWS Mainframe Modernization Competency Partner](#)

AWS Professional Services

AWS Professional Services are a global team of AWS experts that can help you realize your desired business outcomes in the AWS Cloud. To learn more, see [AWS Professional Services](#).

AWS Countdown Premium

[AWS Countdown Premium](#) provides proactive technical guidance during migration environment discovery and scope definition. This guidance prepares customers for planning, testing, cutover and hyper-care. You can also leverage expert AWS knowledge with technical workshops that cover platform and migration capabilities that can upskill customers and institutionalize knowledge. You can accelerate troubleshooting during the Migration Execution Plan and receive expedited troubleshooting and resolution of technical issues with service team engagement.

AWS Countdown Premium can help you migrate with confidence with the help of a designated engineer who can provide context awareness and help you resolve issues, even when it requires additional help from service teams.

AWS Managed Services

[AWS Managed Services \(AMS\)](#) helps you adopt AWS at scale and operate more efficiently and securely. AMS leverages standard AWS services and offers guidance and execution of operational best practices with specialized automations, skills, and experience that are contextual to your environment and applications. To help you focus on innovation, AMS provides proactive, preventative, and detective capabilities that raise the operational bar and help reduce risk without constraining agility. AMS extends your team with operational capabilities including monitoring, incident management, [AWS Incident Detection and Response](#), security, patch, backup, and cost optimization.

Migration tooling

Our comprehensive portfolio of [agentic migration services](#), [migration competency partners](#), and our mature third-party migration tooling ecosystem provide automation and intelligent

recommendations based on AWS machine learning to simplify and accelerate each step of the three-phase migration process.

The following video from re:Invent 2025 describes how to complete a large-scale migration and modernization with the help of agentic AI: [Reimagining large-scale migration planning with agentic AI](#)

See [Education and Enablement](#) for a recommended free course on [Skill Builder](#) that covers the migration process and tooling.

[AWS Transform MGN \(MGN\)](#) provides a solution to migrate business applications from on-premise or other cloud providers easily to AWS.

To learn about the architecture of AWS Transform MGN, watch the following video: [MGN](#)

[AWS Database Migration Service \(AWS DMS\)](#) is a managed migration and replication service that helps move your database and analytics workloads to AWS quickly, securely, and with minimal downtime and zero data loss. AWS DMS supports migration between more than 20 database and analytics engines.

The following video from AWS re:Invent 2023 describes AWS DMS: [Mastering AWS Database Migration Service: Deep Dive into Migration Options, Serverless Capabilities, and Performance Optimization](#)

Case studies

AWS has helped thousands of organizations, including enterprises such as GE, the Coca-Cola Company, BP, Enel, Samsung, NewsCorp, and Twenty-First Century Fox, migrate to the cloud and free-up resources by lowering IT costs while improving productivity, operational resiliency, and business agility.

To explore migration case studies across different industries, workload types, and geographies see the Migration and Modernization [customer stories](#) at the bottom of the page.

Whitepapers and AWS prescriptive guidance

AWS regularly publishes whitepapers that describe various business cases for migrating to AWS together with AWS best practices. To find these whitepapers, see the [whitepaper homepage](#).

[AWS Prescriptive Guidance](#) provides time-tested strategies, guides, and patterns to help accelerate your cloud migration, modernization, and optimization projects. These resources were developed by AWS technology experts and the global community of AWS Partners, based on their years of experience helping customers realize their business objectives on AWS. It contains hundreds of specific migration scenarios. For example, see [Migrate an on-premises Oracle database to Amazon RDS for Oracle](#) and [Export Amazon RDS for SQL Server tables to an S3 bucket by using AWS DMS](#).

AWS Well-Architected Labs

The [AWS Well-Architected Framework](#) describes key concepts, design principles, and architectural best practices for designing and running workloads in the cloud.

These labs help you build and assess AWS infrastructure following AWS architectural and operational best practices.

Watch a [video about the pillars of the AWS Well-Architected Framework](#).

You can also refer to [Migration Lens - AWS Well-Architected Framework](#).

For more information, contact your AWS account team.

Modernization pathways

AWS modernization pathways provide organizations with structured approaches to modernize their application portfolios after migration. Each pathway provides a proven route to specific business outcomes. These include reducing licensing costs, increasing developer velocity, enabling data-driven decisions, and embedding AI into existing workflows.

Organizations rarely follow a single pathway. Most combine multiple pathways across their portfolio based on the specific needs of each workload, team readiness, and business priorities. The pathways that follow are ordered from foundational infrastructure changes to higher-order application transformations, but you can adopt them in any sequence or in parallel.

Move to AI

With this modernization pathway, you can identify and implement high-impact AI opportunities within your existing application portfolio. Rather than building AI initiatives from scratch, this pathway uses current applications as a foundation for AI integration.

When you adopt this pathway, you systematically assess your portfolio for AI opportunities, implement graduated AI capabilities—from basic inference to autonomous agents—and establish the governance frameworks needed for responsible AI adoption.

This pathway provides the following benefits:

- **Unlock trapped value** – Existing applications contain data, workflows, and domain logic that AI can enhance without requiring full rewrites.
- **Accelerate decision-making** – Embed intelligent recommendations, predictions, and automation directly into business workflows where decisions are made.
- **Enhance customer experiences** – Add natural language interfaces, personalization, and predictive capabilities to customer-facing applications.
- **Automate manual processes** – Replace repetitive human judgment tasks (classification, extraction, routing) with AI that improves over time.
- **Competitive differentiation** – Transform commodity applications into intelligent systems that deliver unique value competitors cannot easily replicate.
- **AI-ready architecture** – Modernized applications with clean APIs, proper state management, and observability are ready for safe AI agent integration.

This pathway applies to the following situations:

- Organizations with large application portfolios that contain untapped data and process automation opportunities
- Companies whose competitors are gaining advantage through AI-enhanced experiences
- Teams that have completed initial cloud migration and want to unlock the next wave of business value

Example: A property management company processes thousands of maintenance requests monthly, each requiring manual classification, prioritization, and routing to the appropriate contractor. By integrating AI into the existing request management application, the system automatically classifies incoming requests by urgency, matches them to available contractors based on skill and proximity, and prioritizes them based on tenant impact. Processing time drops from hours to minutes, tenant satisfaction improves, and staff focus shifts from administrative routing to exception handling and quality oversight.

Move to cloud native

Moving to cloud native involves decomposing monolithic applications into smaller, independently deployable services that leverage fully managed cloud capabilities for compute, messaging, storage, and orchestration.

Monolithic architectures have limitations: slow release cycles, cascading failures, and scaling constraints. This pathway addresses each by enabling teams to own discrete business capabilities end-to-end. Applications become more resilient, scalable, and evolvable without requiring coordinated releases across the entire system.

This pathway provides the following benefits:

- **Independent deployability** – Teams ship changes to their service without waiting for or impacting other parts of the system, dramatically increasing release velocity.
- **Granular scaling** – Scale individual components based on demand rather than scaling the entire application, optimizing costs for uneven workload patterns.
- **Fault isolation** – A failure in one service does not cascade to the entire application, improving overall system reliability and customer experience.
- **Technology flexibility** – Each service can use the language, framework, or data store best suited to its problem domain, rather than being constrained by a single technology stack.

- **Faster innovation** – Smaller, focused services are easier to understand, test, and modify, reducing the time from idea to production.

This pathway applies to the following situations:

- Monolithic applications where release cycles are measured in weeks or months
- Systems where one component's failure brings down the entire application
- Organizations where multiple teams are blocked by shared codebases

Example: An insurance company's claims processing system is a single monolith deployed quarterly. Any change—no matter how small—requires a full regression cycle. By decomposing into services aligned to business domains (intake, adjudication, payment, notification), each team deploys independently. The claims intake team ships improvements weekly, while the payment team maintains its own release cadence. A bug in notifications no longer delays claims processing.

Move to containers

Containerization provides a consistent runtime environment that decouples applications from the underlying infrastructure, enabling teams to build, ship, and run software faster and more reliably.

Organizations adopting this pathway package existing applications into containers and deploy them using managed orchestration services. This eliminates environment inconsistencies between development and production, simplifies scaling, and standardizes deployment across teams—regardless of the programming language or framework in use.

This pathway provides the following benefits:

- **Operational consistency** – Containers ensure the same artifact runs identically across development, staging, and production, eliminating the "works on my machine" problem and reducing deployment failures.
- **Resource efficiency** – Multiple containers share the same host, increasing utilization compared to dedicated virtual machines and reducing infrastructure costs.
- **Faster deployments** – Container images start in seconds rather than minutes, enabling rapid scaling and frequent releases without extended maintenance windows.
- **Portability** – Containerized applications are not tied to specific infrastructure, making it straightforward to move workloads across environments or regions.

- **Team autonomy** – Each team can own their deployment pipeline independently, shipping updates without coordinating with other teams or waiting for shared release cycles.

This pathway applies to the following situations:

- Organizations with large VM fleets running diverse applications that want a common deployment standard
- Teams seeking to adopt CI/CD without rewriting applications
- Companies exiting costly virtualization platforms

Example: A retail company operates 200 applications across multiple data centers, each deployed differently depending on when it was built. By containerizing these applications, all teams adopt a single deployment standard. New developers onboard faster, deployments happen multiple times per day instead of monthly, and infrastructure costs drop because workloads share resources more efficiently.

Move to managed databases

Moving to managed databases eliminates the undifferentiated heavy lifting of database administration—patching, backups, replication, failover—and lets teams focus on optimizing data models and queries for their applications.

This pathway covers both migrating from self-managed databases to their managed equivalents and transitioning from commercial engines to open-source alternatives. Organizations gain automated operations, built-in high availability, and pay-as-you-go pricing without sacrificing performance or data integrity.

This pathway provides the following benefits:

- **Reduced operational burden** – Automated patching, backups, point-in-time recovery, and failover eliminate manual DBA tasks that consume engineering time.
- **High availability by default** – Multi-AZ deployments and automated failover provide resilience without complex custom configurations.
- **Cost optimization** – Right-size database instances dynamically, use reserved capacity for predictable workloads, and eliminate over-provisioned hardware.
- **Licensing freedom** – Transitioning from commercial database engines removes per-core or per-socket licensing costs that scale unpredictably with growth.

- **Purpose-built options** – Choose the right database for each workload pattern (relational, document, key-value, graph, time-series) rather than forcing all data into a single engine.

This pathway applies to the following situations:

- Organizations spending significant DBA time on maintenance rather than optimization
- Companies with expiring commercial database licenses
- Teams that need automated disaster recovery and point-in-time restore without custom scripts

Example: A financial services firm runs dozens of self-managed database instances across two data centers, requiring a dedicated team for patching, backup verification, and failover testing. After moving to managed database services, automated backups and cross-region replication run without intervention. The DBA team shifts focus from maintenance to query optimization and data modeling, and the company eliminates annual commercial license renewals.

Move to open source

This pathway helps organizations reduce dependency on proprietary software by migrating to open-source alternatives running on managed cloud infrastructure. It applies to application runtimes, databases, operating systems, and development frameworks.

Beyond cost savings, moving to open source gives organizations transparency into the software they depend on, access to broader talent pools, and freedom from vendor-specific lock-in that constrains architectural decisions.

This pathway provides the following benefits:

- **License cost elimination** – Remove per-seat, per-core, or per-server licensing fees that grow unpredictably as the organization scales.
- **Talent availability** – Open-source skills are more widely available in the market, reducing hiring difficulty and onboarding time for new team members.
- **Community innovation** – Benefit from rapid feature development, security patches, and ecosystem tooling driven by global open-source communities.
- **Architectural freedom** – Make technology decisions based on technical merit rather than license constraints or vendor roadmaps.

- **Transparency and auditability** – Inspect, modify, and audit source code directly, meeting security and compliance requirements without relying on vendor attestations alone.

This pathway applies to the following situations:

- Organizations with significant annual licensing spend on proprietary runtimes or databases
- Companies whose vendor contracts restrict architectural choices
- Teams that struggle to hire for proprietary technology skills

Example: A media company spends millions annually on proprietary runtime licenses across its web platform. Developer hiring is constrained because candidates rarely have experience with the proprietary stack. By migrating to an open-source runtime on managed cloud infrastructure, the company eliminates license costs, attracts from a larger talent pool, and gains access to a rich ecosystem of open-source libraries and tools that accelerate feature development.

Move to modern analytics

This pathway transforms how organizations collect, store, process, and derive insights from data. It replaces fragmented, batch-oriented data pipelines with unified, near-real-time analytics architectures that bring data closer to decision-makers.

Organizations adopting this pathway build modern data foundations that support advanced analytics, machine learning, and business intelligence—breaking down data silos and enabling self-service access across the organization.

This pathway provides the following benefits:

- **Faster time to insight** – Move from batch processing (hours or days) to near-real-time analytics, enabling decisions based on current rather than stale data.
- **Data democratization** – Self-service tools let business users explore data without waiting for engineering teams to build custom reports.
- **Cost-effective storage at scale** – Decouple storage from compute, storing vast amounts of data affordably and scaling processing independently based on query needs.
- **Unified data view** – Consolidate data from disparate sources into a single platform, eliminating inconsistencies between reports from different systems.

- **ML-ready foundations** – Modern data architectures provide the clean, accessible datasets that machine learning initiatives require to succeed.

This pathway applies to the following situations:

- Organizations where reporting takes days due to batch pipelines
- Companies with data locked in silos across business units
- Teams that want to enable self-service analytics without building custom dashboards for every request

Example: A logistics company has shipment data in one system, customer feedback in another, and financial data in a third. Generating a cross-functional report requires manual data extraction by multiple teams. After building a modern analytics foundation, all data flows into a centralized platform accessible to business analysts across the organization. Delivery performance, customer satisfaction, and profitability are visible in a single dashboard updated continuously.

Move to modern DevOps

This pathway modernizes how software is built, tested, and delivered by establishing automated pipelines, infrastructure as code, and continuous feedback loops. It addresses the gap between development speed and operational reliability.

Modern DevOps practices ensure that code changes flow from development to production safely and rapidly, with automated testing, security scanning, and compliance checks built into the pipeline rather than bolted on after the fact.

This pathway provides the following benefits:

- **Accelerated delivery** – Automated pipelines reduce the path from commit to production from weeks to hours or minutes, enabling rapid response to business needs.
- **Shift-left quality** – Automated testing and security scanning catch issues early in the development cycle when they are cheapest to fix.
- **Infrastructure consistency** – Infrastructure as code ensures environments are reproducible, auditable, and drift-free, eliminating configuration-related failures.
- **Reduced toil** – Automation eliminates repetitive manual tasks (deployments, environment provisioning, compliance checks), freeing engineers for higher-value work.

- **Continuous feedback** – Monitoring, alerting, and observability integrated into the delivery pipeline create tight feedback loops between production behavior and development decisions.

This pathway applies to the following situations:

- Organizations where deployments are manual, infrequent, and error-prone
- Teams that lack visibility into production behavior until customers report issues
- Companies where environment provisioning takes days or weeks

Example: A healthcare company deploys software quarterly through a manual checklist involving multiple teams over a weekend maintenance window. Failures during deployment cause extended outages. After adopting modern DevOps practices, deployments happen multiple times per week through automated pipelines. Every change passes through automated tests, security scans, and compliance checks before reaching production. Rollbacks that once took hours now happen in minutes.

For more information, contact your AWS account team.

Generative AI insights for modernization pathways

Once organizations have explored or implemented a [modernization pathway](#), they unlock a new class of opportunities powered by generative AI. Modernized architectures, with clean APIs, event-driven patterns, managed data stores, and containerized workloads, are inherently AI-ready. This page maps the Gen AI insights and use cases that become available as a natural next step along each pathway.

 Note

The modernization + AI flywheel: Agentic AI reduces modernization costs by up to 80% and accelerates transformations 5-80x. Once complete, those modernized workloads become the foundation for Gen AI applications, creating a self-reinforcing cycle where AI enables modernization, and modernization enables more AI.

Source: [Why Agentic AI Marks an Inflection Point for Enterprise Modernization](#), Dr. Asa Kalavade, VP AWS Migration & Modernization

 Note

Key insight: According to McKinsey, generative AI can reduce cloud migration time by 30-40%. But the real value emerges after migration, when modernized applications become the foundation for intelligent automation, agentic workflows, and data-driven decision-making.

Source: [Accelerating Cloud Migration with AWS Transform and Generative AI](#)

How modernization enables Gen AI

Modernization outcome	Gen AI enabler
Clean, well-documented APIs	AI agents can safely invoke application functions as tools
Event-driven architectures	Real-time triggers for AI inference and autonomous actions

Modernization outcome	Gen AI enabler
Managed databases with structured data	RAG knowledge bases, semantic search, and analytics
Containerized microservices	Independent scaling of AI inference alongside business logic
Observability and logging	AI-powered root cause analysis and predictive operations
Open-source frameworks	Flexibility to integrate open-weight models and custom fine-tuning

Gen AI insights by pathway

Move to AI: Agentic AI and autonomous workflows

Prerequisite modernization: Applications with modular APIs, proper state management, and observability.

Once your portfolio is modernized, you can implement graduated AI capabilities, from basic inference to fully autonomous agents:

AI maturity level	Capability	Example
Level 1: Inference	Single-turn predictions and classifications	Auto-classify support tickets by urgency
Level 2: RAG-enhanced	Context-aware generation using enterprise knowledge	Maintenance request routing with contractor matching
Level 3: Agents	Multi-step reasoning with tool use	Customer onboarding agent that checks eligibility, creates accounts, and schedules follow-ups

AI maturity level	Capability	Example
Level 4: Multi-agent	Orchestrated teams of specialized agents	Supply chain optimization with demand forecasting, inventory, and logistics agents working in concert

Insight from the field: "By 2028, over 33% of enterprise applications will embed agentic capabilities, up from less than 1% today." (Gartner, cited in [Effectively Building AI Agents on AWS Serverless](#))

Move to cloud native: Intelligent microservices and event-driven AI

Prerequisite modernization: Decomposed services with independent deployment, event buses, and API gateways.

Cloud-native architectures are the ideal substrate for embedding AI into specific business domains without monolithic AI projects:

Gen AI opportunities unlocked:

- **Per-service intelligence** - Each microservice can embed its own AI capability (NLU for customer service, anomaly detection for payments, recommendations for product catalog) without coupling AI decisions across the system. (Reference: [Drive hyper-personalized customer experiences with Amazon Personalize and generative AI](#))
- **Event-driven inference** - Trigger AI workflows automatically when business events occur (order placed, fraud check, personalized confirmation). (Reference: [Effectively Building AI Agents on AWS Serverless](#))
- **Serverless AI agents** - Deploy autonomous agents that scale to zero when idle and burst on demand. No GPUs required for orchestration. (Reference: [Building Serverless Architectures for Agentic AI on AWS](#))
- **Real-time personalization** - Decoupled services can each contribute context to a central personalization engine, enabling hyper-personalized experiences. (Reference: [Drive hyper-personalized customer experiences with Amazon Personalize and generative AI](#))

Move to containers: Scalable AI inference and model serving

Prerequisite modernization: Containerized workloads with CI/CD pipelines and managed orchestration.

Containerization creates a consistent packaging model that applies equally well to AI workloads:

Gen AI opportunities unlocked:

- **Sidecar AI patterns** - Deploy lightweight inference containers alongside application containers for real-time enrichment (for example, content moderation, sentiment analysis, translation). (Reference: [Optimizing a lift-and-shift for security](#))
- **Custom model hosting** - Run fine-tuned or open-weight models (Llama, Mistral) alongside your application tier for data-sensitive use cases.
- **GPU-accelerated inference** - Use GPU node pools for latency-sensitive generative tasks (image generation, real-time transcription).
- **A/B testing AI models** - Container orchestration enables canary deployments of different model versions, measuring business impact before full rollout.

Move to managed databases: RAG, semantic search, and data intelligence

Prerequisite modernization: Data in managed database services with automated operations.

Managed databases are the fuel for generative AI. They provide the structured and unstructured data that makes AI responses accurate and contextual:

Gen AI opportunities unlocked:

- **Enterprise RAG (Retrieval-Augmented Generation)** - Connect knowledge bases directly to your data stores. The system automatically ingests, chunks, embeds, and indexes your data for semantic retrieval. (Reference: [Build generative AI applications with Amazon Aurora and Amazon Bedrock Knowledge Bases](#))
- **Natural language queries** - Allow business users to ask questions of databases in plain English. AI translates natural language to SQL and returns formatted answers. (Reference: [Query Amazon Aurora PostgreSQL using Amazon Bedrock Knowledge Bases structured data](#))
- **Automated data quality** - Use generative AI to detect anomalies, fill gaps, standardize formats, and generate data documentation automatically.

- **Knowledge graph enrichment** - Extract entities and relationships from unstructured data stored in your databases to build organizational knowledge graphs.

Reference: [How Amazon Bedrock Knowledge Bases work](#)

Move to open source: Flexible model selection and fine-tuning

Prerequisite modernization: Applications running on open-source frameworks, databases, and runtimes.

Open-source modernization extends naturally into the AI layer, giving teams maximum control over model selection and customization:

Gen AI opportunities unlocked:

- **Open-weight model deployment** - Host models like Meta Llama, Mistral, or DeepSeek for full control over inference without vendor lock-in.
- **Domain-specific fine-tuning** - Fine-tune open models on your proprietary data for industry-specific accuracy (legal, medical, financial terminology). (Reference: [Automate fine-tuning of Llama 3.x models with Amazon SageMaker Pipelines](#))
- **Cost optimization** - Use smaller, specialized open models for routine tasks (classification, extraction) and reserve larger models for complex generation tasks.
- **Community-driven innovation** - Leverage open-source AI frameworks (LangChain, Strands Agents SDK, LlamaIndex) that integrate seamlessly with open-source application stacks. (Reference: [Effectively Building AI Agents on AWS Serverless](#))

Move to modern analytics: Generative BI and predictive insights

Prerequisite modernization: Analytics workloads on modern data platforms with clean data pipelines.

Modern analytics architectures create the data foundation for generative BI, where AI doesn't just report on the past but predicts the future and recommends actions:

Gen AI opportunities unlocked:

- **Natural language analytics** - Business users ask questions in plain language and receive charts, summaries, and insights without writing SQL or building dashboards. (Reference: [Pioneering Generative AI for Business Intelligence with Amazon Q in QuickSight](#))
- **Automated narrative generation** - Generate executive summaries, trend explanations, and anomaly reports automatically from analytics outputs.
- **Predictive recommendations** - Move from descriptive ("what happened") to prescriptive ("what should we do") by combining analytics data with LLM reasoning.
- **Data storytelling** - AI generates contextual narratives around data changes, helping non-technical stakeholders understand trends and make decisions.

Move to modern DevOps: AI-powered operations and developer productivity

Prerequisite modernization: CI/CD pipelines, infrastructure as code, observability, and automated testing.

Modern DevOps practices create the automation backbone that generative AI can supercharge:

Gen AI opportunities unlocked:

- **Autonomous incident investigation (AWS DevOps Agent)** - A frontier AI agent that autonomously investigates production issues the moment they occur. It correlates telemetry, code, and deployment data to identify root cause, completing investigations in under 5 minutes without human intervention. (Reference: [Leverage Agentic AI for Autonomous Incident Response with AWS DevOps Agent](#))
- **Code generation and review** - AI-powered development assistants accelerate development with context-aware code suggestions, test generation, and security scanning.
- **Proactive incident prevention** - AWS DevOps Agent analyzes patterns across historical incidents to deliver targeted improvements across observability, infrastructure optimization, deployment pipelines, and application resilience.
- **Release readiness review** - AI reviews code changes during generation, checking policy compliance, dependency impacts, blast radius, and access controls before merge.
- **Autonomous release testing** - Generates and runs change-specific test plans targeting risk areas, catching regressions and integration failures before production.
- **Shift-left security** - Generative AI scans code, infrastructure templates, and configurations for vulnerabilities during development rather than after deployment.

AWS DevOps Agent, key capabilities:

Area	Capability	How it works
Production operations	Automated incident investigation	Begins investigating immediately from alerts/tickets; correlates telemetry, code, and deployments to identify root cause
Production operations	Actionable mitigation plans	Provides specific actions to resolve incidents; generates agent-ready instructions for implementation by tools like Kiro
Production operations	Proactive prevention	Analyzes incident patterns to recommend observability, infra, pipeline, and resilience improvements
Production operations	Incident coordination	Routes findings through Slack, ServiceNow, PagerDuty ; creates support cases with full investigation context
Production operations	On-demand SRE tasks	Handles SRE queries via natural language: resource health, incident patterns, deployment tracking
Release management (preview)	Release readiness review	Reviews code changes for policy compliance, cross-repo dependencies, blast radius, and permission escalation
Release management (preview)	Autonomous release testing	Generates and runs change-specific test plans in customer

Area	Capability	How it works
		environments, targeting risk areas surfaced in review
Release management (preview)	Developer workflow integration	Delivers results through PRs, coding agent IDEs, and CI/CD pipelines

Real-world example: A monitoring alarm triggers due to elevated 5xx errors. AWS DevOps Agent autonomously tests hypotheses, identifies database write throttling from a recent deployment, and posts root cause analysis with mitigation recommendations to Slack, all in under 5 minutes from alarm to actionable solution. (Reference: [Resolve Application Issues Autonomously with AWS DevOps Agent and Dynatrace](#))

References:

- [Leverage Agentic AI for Autonomous Incident Response with AWS DevOps Agent](#)
- [Resolve Application Issues Autonomously with AWS DevOps Agent and Dynatrace](#)
- [Generative AI-powered Technology Operations](#)

Summary

Modernization is not the end state. It is the launchpad for generative AI. Each pathway creates specific architectural preconditions that make AI integration natural, safe, and cost-effective:

Pathway	AI unlock
Move to AI	Agentic workflows and autonomous agents
Move to cloud native	Event-driven, per-service intelligence
Move to containers	Scalable model serving and sidecar inference
Move to managed databases	Enterprise RAG and natural language data access

Pathway	AI unlock
Move to open source	Flexible model selection and domain fine-tuning
Move to modern analytics	Generative BI and predictive insights
Move to modern DevOps	AI-powered operations and self-healing systems

Organizations that have completed modernization are uniquely positioned to embed AI into their operations, not as a separate initiative, but as a natural extension of the architecture they have already built.

For more information, contact your AWS account team.

Organizational readiness

Organizational readiness is a critical factor for maximizing value in the cloud. Organizations that assess their gaps early and put in place the right transformation plans have significantly higher chances of achieving their business goals when moving to the cloud.

Some of the most common blockers for cloud value realization for customers include:

- Lack of visible and active sponsorship
- Siloed workflows between organizational units
- Architectural entanglement
- Undefined operating model
- Analysis paralysis
- Talent and skills gaps
- Misaligned teams
- Unrealistic goals

At AWS, we have helped thousands of organizations in their cloud adoption and have captured our experience and best practices through different frameworks and tools.

In the following section, we introduce how you can organize your key people around cloud adoption through the Cloud Center of Excellence.

Then, we cover the Cloud Adoption Framework (CAF), which covers the 6 key organizational capabilities that underpin successful cloud transformations, as well as the Cloud Adoption Readiness Tool, a self-service assessment for your organization to measure alignment against the six perspectives of the CAF.

Finally, in the Cloud Operations section, we go through some resources to ensure your organization can operate your workloads in the cloud securely and reliably.

Cloud Center of Excellence (CCoE)

The Cloud Center of Excellence (CCoE), also known as Cloud Enablement Engine (CEE), is a multi-disciplinary team in charge of implementing the governance, best practices, training, and architecture needed for cloud adoption in a manner that provides repeatable patterns for the

larger enterprise to follow. The CCoE team is made up of the Cloud Business Office (CBO) and the Cloud Platform Engineering (CPE), and is governed by the Cloud executive team. The team can be global, local, or federated (though some [capabilities](#) may be local and others global). In a migration context, the CCoE is typically started in the [Mobilize Phase](#). For more resources, see [Accelerating cloud adoption through culture, change, and leadership](#). You can also watch the video, [The human side of cloud transformation | Amazon Web Services](#). For more information, see [example CCoE structure](#) and [prescriptive guidance on building CCoEs](#).

AWS Cloud Adoption Framework (CAF)

The [AWS Cloud Adoption Framework \(AWS CAF\)](#) leverages AWS experience and best practices to help you digitally transform and accelerate your business outcomes through innovative use of AWS. AWS CAF identifies specific organizational capabilities that underpin successful cloud transformations. These capabilities provide best-practice guidance that helps you improve your cloud readiness. AWS CAF groups its capabilities in six perspectives: Business, People, Governance, Platform, Security, and Operations. Each perspective comprises a set of capabilities that functionally related stakeholders own or manage in the cloud transformation journey. Use the AWS CAF to identify and prioritize transformation opportunities, evaluate and improve your cloud readiness, and iteratively evolve your transformation roadmap.

Cloud Readiness Assessment (CRA)

No matter which stage your organization is in on the cloud journey, it is never too late to assess your cloud readiness. The [AWS Cloud Readiness Assessment \(CRA\)](#) is a free self-service online assessment tool to evaluate your cloud readiness and identify gaps and opportunities against the AWS CAF capabilities, across six perspectives. The outcome is a detailed report with recommended actions to help your organization achieve the maximum benefits of cloud adoption.

Cloud Maturity Assessment (CMA)

The AWS CMA is a maturity model that tracks the adoption of a customer's cloud transformation over time. It uses AWS CAF's foundational organizational capabilities that underpin an organization's ability to leverage the cloud to digitally transform in the cloud. It is an assessment led by AWS with the inputs from your organization to identify risks and create an engagement and risk mitigation plan in support of achieving customer goals. The results are tracked over time, and can be benchmarked, and thus used as an important input for decision making.

Cloud Operations

Operating your applications in the cloud is different from operating them on-premises (whether in your own data centers or in co-location). In your cloud migration journey, your organization will gradually pivot towards a Cloud Operating Model. The [Well-Architected Framework - Operational Excellence Pillar](#) can be very useful when designing your cloud operating model based on your Application Engineering and Operations (AEO) and Infrastructure Engineering Operations (AEO) teams. To learn more about Cloud Operating Models, watch [AWS re:Invent 2024 - Building the future of cloud operations at any scale](#).

For additional guidance around capabilities needed to operate in the cloud, see [AWS Cloud Adoption Framework: Operations Perspective](#).

For prescriptive guidance on how to build your own Cloud Operating Model according to the AWS definition, see [Building your Cloud Operating Model](#). For tips on how to make sure the platform that you build is effective in supporting your application development teams, see [How to Set Up a Platform That Effectively Supports Your Development Teams](#).

Maturity of cloud operations

You can assess and improve the maturity of your cloud operations with the help of the [Well-Architected Tool](#). The tool makes it easy for your teams to review a given application against the best practices of the Operational Excellence and Security pillars of the AWS Well-Architected Framework.

Cloud Operations can be improved with [AWS GameDays](#), [AWS Workshops](#), [Table Top simulations](#) and [Well-Architected Labs](#).

Additional resources

The following two-part blog provides the foundational tooling that will help you centralize and automate operations, and improve governance and visibility through AWS Cloud Operations services: [Increase visibility and governance on cloud with AWS Cloud Operations services – part 1](#) and [part 2](#).

Cloud financial management and FinOps

Understanding the processes around how resources are priced and consumed in the cloud is one of the critical capabilities that need to be developed during your migration. [AWS Cloud Financial](#)

[Management](#) can help you find initial resources and guidance on how to start and evolve your [Cloud Financial Management capability](#) and FinOps. To dive deeper into this topic, watch the following AWS re:Invent video that discusses the practice: [FinOps 3.0: Cost Intelligence for the AI Era](#).

For key 2025 re:Invent launches for CFM, see [AWS Cloud Financial Management key re:Invent 2025 launches](#).

For a customer case study, see [Petabyte-Scale Cost Optimization: How a Video Hosting platform Saved 70% on S3](#)

For more information, contact your AWS account team.

Operational readiness

A successful approach to operations is critical to the success of every organization and its digital transformation. [Operational excellence](#) is required to ensure that your transformation achieves its purpose and that applications consistently meet their business outcomes and the expectations of their users. This topic is intended to help you get a head start on your operations.

Managing your resources

Use an orchestration tool to manage, configure, patch, and use automation.

AWS Systems Manager: Use [AWS Systems Manager Quick Setup](#) to quickly configure required security roles and commonly used Systems Manager capabilities on your Amazon Elastic Compute Cloud instances across your entire organization.

The following are some of the capabilities of Systems Manager:

- [AWS Systems Manager Session Manager](#) allows you to SSH to Linux or RDP to Windows instances without the need to open inbound ports, maintain bastion hosts, or manage SSH keys.
- [AWS Systems Manager Distributor](#) allows you to package and publish your own software or to find and publish AWS-provided agent software packages, such as AmazonCloudWatchAgent, or third-party packages such as Trend Micro.
- [AWS Systems Manager Patch Manager](#) allows you to automate the process of patching managed nodes with both security-related updates and other types of updates.

Monitoring and observability

Implement observability in infrastructure and applications so that you can understand their state and make data-driven decisions based on business requirements.

- **Amazon CloudWatch Metrics:** Centralize performance data with basic monitoring metrics at no charge from AWS services such as Amazon EC2, Amazon Elastic Block Store (Amazon EBS), and Amazon Relational Database Service (Amazon RDS). For more information, see [AWS services that publish CloudWatch metrics](#) in the Amazon CloudWatch User Guide.
- **CloudWatch Dashboards:** To learn how to see key metrics from all AWS services, refer to [See the pre-built cross-service dashboard](#). To learn how to use automatic dashboards to focus on

metrics and alarms in a single AWS service, refer to [See a pre-built dashboard for a single AWS service](#). The metrics that are considered most important by service teams are used on automatic dashboards. Use these metrics as a starting point for your own dashboards.

- **CloudWatch Logs:** To learn how to centralize, manage, and analyze logs from AWS services, see [Enabling logging from AWS services](#) in the Amazon CloudWatch Logs User Guide.
- **CloudWatch Alarms:** A metric alarm watches a single CloudWatch metric or the result of a math expression based on CloudWatch metrics. The alarm performs one or more actions based on the value of the metric or expression relative to a threshold over a number of time periods. To learn how to create and use alarms, see [Using Amazon CloudWatch alarms](#).
- **CloudWatch Agent:** [Install the CloudWatch agent](#) as part of [AWS Systems Manager Quick Setup](#) to send instance metrics and logs to CloudWatch. Use these metrics and logs alongside AWS service metrics and logs to create dashboards and alarms for your infrastructure and applications.
- **AWS CloudTrail:** CloudTrail is active in your AWS account when you create it and doesn't require any manual setup. When activity occurs in your AWS account, that activity is recorded in a CloudTrail event. [Tutorial: Review AWS account activity in event history](#)

Set up a monitoring account for cross-account observability. To learn how, see [Link monitoring accounts with source accounts](#) in the Amazon CloudWatch User Guide.

Security and compliance

Ensuring the security and compliance of your AWS environment is paramount.

- **AWS Security Hub CSPM:** Security Hub provides you with a comprehensive view of your security state in AWS and helps you assess your AWS environment against security industry standards and best practices. To learn how to set it up and use it, see [Enabling Security Hub](#) in the AWS Security Hub User Guide.
- **AWS Config:** AWS Config provides a detailed view of the configuration of AWS resources in your AWS account. This includes how the resources are related to one another and how they were configured in the past so that you can see how the configurations and relationships change over time. Use AWS Config to assess, audit, and evaluate the configurations of your AWS resources. To learn how to set it up, see [Getting Started with AWS Config](#) in the AWS Config Developer Guide.

Video: [Back to Basics: Using AWS Config; and Conformance Packs to Optimize Your AWS Resources](#)

Backup

AWS Backup is a fully-managed service that helps you centralize and automate data protection across AWS services, in the cloud, and on premises. To learn more, see [Getting started with AWS Backup](#) in the AWS Backup Developer Guide.

Optimization and cost management

Optimize your AWS usage to ensure cost efficiency without compromising on performance and security.

- **AWS Cost Explorer:** This is a tool that you can use to visualize, understand, and manage your AWS costs and usage over time. To learn how to enable it and use it, see [Analyzing your costs with AWS Cost Explorer](#) in the AWS Cost Management User Guide.
- **AWS Trusted Advisor:** Trusted Advisor draws upon best practices learned from helping hundreds of thousands of AWS customers. It inspects your AWS environment, and then makes recommendations when opportunities exist to save money, improve system availability and performance, or help close security gaps. For more information, see [AWS Trusted Advisor](#) in the AWS Support User Guide.
- **AWS Compute Optimizer:** This is a service that analyzes the configuration of your AWS resources and their utilization metrics to provide you with right-sizing recommendations. Use Compute Optimizer to right-size workloads according to your workload preferences through artificial intelligence and machine-learning-based analytics to reduce costs by up to 25%. By using [memory metrics](#) collected by the CloudWatch agent, the recommendations for right-sizing EC2 instances are improved. For more information, see [Getting started with AWS Compute Optimizer](#).

For more information, contact your AWS account team.

Education and enablement

Education and enablement play a fundamental role in cloud transformation by ensuring that employees have the knowledge, skills, and mindset needed to adapt to change, contribute effectively to the transformation process, and ultimately help the organization achieve its strategic goals.

This page contains the most commonly used mechanisms to enable your organization to fully leverage AWS tools and services.

Learning-needs analysis

You can identify your organization's cloud skills gaps with our free self-assessment tool, the [AWS Learning Needs Analysis](#).

AWS training and certification

[AWS Training and Certification](#) validates cloud expertise to help professionals highlight in-demand skills and organizations build effective, innovative teams for cloud initiatives using AWS. Choose from diverse certification exams by role and specialty designed to empower individuals and teams to meet their unique goals.

Watch the video [AWS Training and Certification – Building AWS Cloud Skills and Confidence](#).

AWS Skill Builder

[AWS Skill Builder](#) is the go-to place for AWS training and certification. You can find learning paths by role or technology area, and browse playlists of videos, tutorials, and documents to improve your AWS skills. Skill Builder currently has over 600 courses at no charge. With AWS Skill Builder individual or team subscriptions, you also gain access to **AWS Builder Labs** that provide hands-on guided exercises to help you develop practical migration skills.

Watch a [video about AWS Skill Builder](#).

A recommended starting point for preparing your organization technically to migrate to AWS is the [Migration Foundations](#) learning path. It comes with learning plans, ramp-up guides, assessment,

and the following digital badge that you can include in your résumé or share with your employer or on social media after you pass the assessment.



AWS GameDay

GameDay is a collaborative learning exercise that tests skills in implementing AWS solutions to solve real-world problems in a gamified, risk-free environment. This is a completely hands-on opportunity for technical professionals to explore AWS services, architecture patterns, best practices, and group cooperation.

Browse the [AWS GameDay](#) website for upcoming public GameDays. It is a hands-on technical training delivered in a gamified and competitive way. It can be used to safely try out new services and fosters collaborative problem-solving.

Watch the video [From Marine Science to Cloud Engineering: How AWS Certification Transformed a Career | AWS Success Story](#).

AWS Workshops

[Workshops](#) are hands-on events designed to teach or introduce practical skills, techniques, or concepts which you can use to solve business problems.

Watch the following video: [Accelerating Learning with AWS Workshops](#).

Experience-Based Acceleration

[Experience-Based Acceleration \(EBA\)](#) is a hands-on and agile-approach to execute specific cloud objectives. We use this in workshops called 'EBA Parties', aimed at addressing friction points and blockers to an organization's cloud journey. In these parties, customer and partner's builders develop new ways of working and collaboration. Together they solve problems while learning best practices in an interactive and immersive setting. Reach out to your AWS account team to discuss your EBA options.

For an introduction to EBA, watch the following video: [AWS re:Invent 2024 - Lockheed Martin's path to digital transformation through innovation](#).

For more information, contact your AWS account team.

Getting support

[AWS Support](#) provides a mix of tools and technology, people, and programs designed to proactively help you optimize performance, lower costs, and innovate faster. We save time for your team by helping you move faster in the cloud and focus on your core business. We are determined to make our customers successful on their cloud journey and address requests that range from answering best practices questions, guidance on configuration, all the way to break-fix and problem resolution.

To learn how to create a support case and manage your open cases with AWS, see [Creating support cases and case management](#) in the AWS Support User Guide. The guide explains the process and also a specific example for a question about a customer's account and billing.

[AWS Enterprise Support](#) provides you concierge-like service where the main focus is helping you achieve your outcomes and find success in the cloud. With Enterprise Support, you get 24/7 technical support from high-quality engineers, tools and technology to automatically manage the health of your environment. You also get access to a Technical Account Manager (TAM) who will provide consultative architectural and operational guidance delivered in the context of your applications and use-cases to help you achieve the greatest value from AWS.

For more information, contact your AWS account team.

Resources

This page aggregates links referenced throughout the Migration Launch Guide for quicker access.

Migration Acceleration Program (MAP)

- [MAP homepage](#)
- Introductory video explaining the MAP [\[English\]](#) [\[French\]](#) [\[Italian\]](#) [\[Spanish\]](#)
- [AWS Migration Acceleration Program 2.0 Tagging Guide](#)
- [List of MAP credit eligible services](#)
- Video on how to get started with tagging for MAP [\[English\]](#) [\[French\]](#) [\[Italian\]](#) [\[Spanish\]](#)
- [AWS Documentation: Tagging AWS Resources](#)
- [Blog Post: Maximizing resource tagging at scale and across teams for your migration to AWS](#)
- Training video explaining how to tag resources on AWS [\[English\]](#) [\[Spanish\]](#)
- [Tagging with CloudFormation](#)
- [AWS Config Conformance Pack](#)
- [Service control policies \(SCPs\)](#)
- [Enforcing tagging on resource creation](#)
- [Tag Policies user guide](#)
- [Restricting allowed tag key/values](#)
- [Amazon Bedrock MAP Tagging using the AWS CLI](#)
- [User guide for MAP 2.0 customer facing dashboard](#)

Delivering a migration program

- [AWS cloud migration website](#)
- [Mobilize your organization to accelerate large-scale migrations](#)
- [Large migration guidelines](#)
- [MAP Migration Process](#)
- [Application Portfolio Assessment Strategy](#)
- [Portfolio assessment guide](#)

- [AWS migration process and its phases](#)
- [Video about the Migration Evaluator](#)
- [Video about the Mobilize phase](#)
- [AWS Migration Hub Journeys](#)
- [MAP Overview | Amazon Web Services \(video\)](#)
- [Build a cloud center of excellence | AWS Events \(video\)](#)
- [AWS Competency Partner Program](#)
- [Managed Service Providers \(MSP\)](#)
- [AWS Partners for migration](#)
- [AWS Partners for Windows](#)
- [AWS Partners for SAP](#)
- [AWS Partners for Oracle](#)
- [AWS Partners for Mainframe](#)
- [AWS Professional Services \(ProServe\)](#)
- [AWS Countdown](#)
- [AWS Managed Services \(AMS\)](#)
- [AWS Incident Detection and Response](#)
- [AWS Migration Hub](#)
- [Third-party migration tooling ecosystem](#)
- [Reimagining large-scale migration planning with agentic AI \(video\)](#)
- [AWS Transform MGN \(MGN\)](#)
- [Video about AWS Transform MGN](#)
- [Migration Foundations - Knowledge Badge](#)
- [AWS Database Migration Service \(AWS DMS\)](#)
- [Mastering AWS Database Migration Service \(video\)](#)
- [AWS Prescriptive Guidance Publications](#)
- [Migration and Modernization customer stories](#)
- [AWS whitepaper homepage](#)
- [AWS Prescriptive Guidance portal](#)
- [AWS Well-Architected Framework](#)

- [Well-Architected Labs](#)
- [Migration Lens - AWS Well-Architected Framework](#)
- [AWS Transform MGN workshop repository](#)
- [AWS Transform](#)
- [AWS Transform assessments](#)
- [Video about the pillars of the AWS Well-Architected Framework](#)

Organizational readiness

- [AWS Cloud Adoption Framework](#)
- [Full documentation about CAF](#)
- [AWS Cloud Adoption Readiness Tool \(CART\)](#)
- [Cloud Enablement Engine \(CEE\)](#)
- [CCoE capabilities](#)
- [People Perspective](#)
- [Role of CCoE in Mobilize Phase](#)
- [The human side of cloud transformation | Amazon Web Services \(video\)](#)
- [Cloud Business Office \(CBO\)](#)
- [Cloud Operating Models](#)
- [How to organize teams to operate a containers platform](#)
- [AWS re:Invent 2024 - Building the future of cloud operations at any scale \(video\)](#)
- [CAF Operations perspective](#)
- [Well-Architected Framework - Operational Excellence Pillar](#)
- [2x2 diagrams to identify source operating models](#)
- [Prescriptive guidance for creating your Cloud Operating Model](#)
- [Tips to make your platform team effective in supporting your application development teams](#)
- [Well-Architected Tool](#)
- [Tabletop simulations](#)
- [Well-Architected Labs on operational excellence](#)
- [re:Invent 2022 - Transform with AWS Cloud Operations: From vision to reality](#)
- [Increase visibility and governance on cloud with AWS Cloud Operations services – part 1](#)

- [Increase visibility and governance on cloud with AWS Cloud Operations services – part 2](#)
- [AWS Cloud Financial Management](#)
- [Cloud Financial Management capability](#)
- [Cloud Value Framework](#)
- [FinOps 3.0: Cost Intelligence for the AI Era \(video\)](#)
- [Key 2025 re:Invent Launches for CFM](#)
- [Example CCoE structure](#)
- [Prescriptive guidance on building CCoEs](#)
- [Petabyte-Scale Cost Optimization: How a Video Hosting platform Saved 70% on S3](#)

Operational readiness

- [AWS Cloud Adoption Framework: Operations Perspective](#)
- [Operational Excellence Pillar - AWS Well-Architected Framework](#)
- [AWS Whitepaper - Change Management in the Cloud](#)
- [Management and Governance on AWS](#)
 - [AWS Systems Manager](#)
 - [Amazon CloudWatch](#)
 - [AWS CloudTrail](#)
 - [AWS Security Hub CSPM](#)
 - [AWS Config](#)
 - [AWS Backup](#)
 - [AWS Cost Explorer](#)
 - [AWS Trusted Advisor](#)
 - [AWS Compute Optimizer](#)
- [Cloud Financial Management with AWS](#)
- [AWS Cloud Operations & Migrations Blog](#)
- [AWS Security Blog](#)
- [AWS Skill Builder - Management & Governance](#)
- [AWS Skill Builder - Security, identity, and compliance](#)
- [AWS Back to basics video playlist](#)

- [AWS Workshops](#)
 - [One Observability Workshop](#)
 - [Security Hub Workshop](#)
 - [Operational excellence through automated Trusted Advisor remediations Workshop](#)

Education and enablement

- [AWS Learning Needs Analysis](#)
- [AWS Training and Certification](#)
- [Video about AWS Training and Certification](#)
- [AWS Skill Builder](#)
- [Video about AWS Skill Builder](#)
- [AWS GameDay](#)
- [Video about AWS GameDay](#)
- [AWS Workshops](#)
- [Video about Accelerating Learning with AWS Workshops](#)
- [Blog - Level up your Cloud Transformation with Experience-Based Acceleration \(EBA\)](#)
- [AWS re:Invent 2024 - Lockheed Martin's path to digital transformation through innovation \(video\)](#)
- [From Marine Science to Cloud Engineering: How AWS Certification Transformed a Career | AWS Success Story \(video\)](#)

Generative AI and modernization

- [Why Agentic AI Marks an Inflection Point for Enterprise Modernization](#)
- [Effectively Building AI Agents on AWS Serverless](#)
- [Building Serverless Architectures for Agentic AI on AWS](#)
- [Generative AI-Powered Technology Operations](#)
- [How Amazon Bedrock Knowledge Bases Work](#)
- [Accelerating Cloud Migration with AWS Transform and Generative AI](#)
- [Leverage Agentic AI for Autonomous Incident Response with AWS DevOps Agent](#)

- [Resolve Application Issues Autonomously with AWS DevOps Agent and Dynatrace](#)
- [Drive hyper-personalized customer experiences with Amazon Personalize and generative AI](#)
- [Build generative AI applications with Amazon Aurora and Amazon Bedrock Knowledge Bases](#)
- [Query Amazon Aurora PostgreSQL using Amazon Bedrock Knowledge Bases structured data](#)
- [Automate fine-tuning of Llama 3.x models with Amazon SageMaker Pipelines](#)
- [Generative AI capabilities for operations and maintenance](#)
- [Optimizing a lift-and-shift for security](#)
- [Pioneering Generative AI for Business Intelligence with Amazon Q in QuickSight](#)
- [Agentic AI vs Generative AI](#)

For more information, contact your AWS account team.