
AWS Prescriptive Guidance

Migrating SAS Grid to the AWS Cloud



AWS Prescriptive Guidance: Migrating SAS Grid to the AWS Cloud

Copyright © 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

Home	1
SAS Grid migration at a glance	1
Assumptions and prerequisites	3
Migration requirements	3
Knowledge requirements	3
Additional SAS considerations	3
Cost and licensing	5
SAS licenses	5
Cost estimates for AWS services	6
Architecture	7
SAS infrastructure	7
Considerations for SAS on AWS	9
SAS background	9
SAS shared file system	9
SAS Grid server tier instance types	9
SAS Grid middle tier and metadata server tier instance types	10
High availability and disaster recovery for SAS Grid	10
Target architecture	10
Automation and tooling	13
High-level migration steps	15
Roles and responsibilities	16
Prepare – Discovery and assessment	17
Prepare current SAS Grid workload assessment	17
Prepare SAS security assessment	17
Prepare SAS Grid migration assessment	17
Set up a new AWS environment (for new AWS users only)	18
Deploy – SAS software on AWS	19
Migrate – SAS content to AWS	20
Migrate Active Directory identities	20
Migrate Linux files, directories, and permissions	20
Migrate SAS metadata	20
Perform post-migration validation and acceptance testing	20
Migrating data associated with SAS	22
Onboarding new customers	23
Additional resources	24
AWS Prescriptive Guidance glossary	25
Document history	32

Migrating SAS Grid to the AWS Cloud

Battulga Purevragcha, Partner Solutions Architect, Amazon Web Services

Eric Y. Yu, Solutions Architect, SAS Institute

July 2020

This guide provides prescriptive steps to streamline the migration of SAS Grid software to Amazon Web Services (AWS).

SAS customers migrate their applications from on-premises data centers to AWS to gain access to cloud-based data lakes and data warehouses. Their goals are to increase the agility, security, and reliability of their applications, to lower costs, and to improve data analytics capabilities. Moving a SAS software deployment to a new location is a multi-step process that involves tasks that must be identified, planned, implemented, and tested.

This guide is intended for organizations that want to rehost or replatform their current SAS Grid installations from on premises or privately hosted environments to AWS. This migration enables organizations to evolve analytics capabilities, minimize rehosting or replatforming risks, and standardize governance and management of the statistical computing environment on AWS. The target audience is IT professionals who have both SAS and AWS expertise.

SAS Grid migration at a glance

Workload	Source workload	<ul style="list-style-type: none">• SAS Grid Manager for Platform• SAS Grid Manager
	Source environment	<ul style="list-style-type: none">• Unix, Linux• On-premises/co-location/non-AWS environment
	Destination workload	<ul style="list-style-type: none">• SAS Grid Manager• Deployment: SAS Intelligence Platform software on a multi-machine host
	Destination environment	<ul style="list-style-type: none">• AWS• Operating model: customer/MSP (ISV)
Migration	Migration strategy (7 Rs (p. 25))	Rehost/replatform
	Is this an upgrade in workload version?	No

AWS Prescriptive Guidance Migrating
SAS Grid to the AWS Cloud
SAS Grid migration at a glance

	Is the source workload different from the ISV workload?	No
	Migration duration	Varies by customer
Cost	Cost of running ISV workload on AWS	Cost and licensing (p. 5)
	Cost of running ISV associated workload that is being migrated to AWS	No
Assumptions and prerequisites	System limitations (minimum/maximum requirements)	SAS System Requirements
	Service-level agreements (SLAs)	SAS Technical Support Services and Policies
	Recovery time objective (RTO)	SAS 9.4 Disaster Recovery Policy
	Recovery point objective (RPO)	SAS 9.4 Disaster Recovery Policy
	Licensing and operating model for the target AWS account	<ul style="list-style-type: none"> • Bring Your Own License (BYOL) • Managed services
	Migration tooling	<ul style="list-style-type: none"> • SAS Migration Utility • AWS Database Migration Service (AWS DMS)
	AWS services used	<ul style="list-style-type: none"> • Amazon Elastic Compute Cloud (Amazon EC2) • Amazon FSx for Lustre • Amazon Simple Storage Service (Amazon S3) • Amazon Virtual Private Cloud (Amazon VPC) <ul style="list-style-type: none"> • NAT gateway • Internet gateway • Amazon EC2 Auto Scaling • AWS Identity and Access Management (IAM)
	Benchmarks	Contact the SAS Enterprise Excellence Center for benchmark information relevant to your site.
Compliance	Security and compliance requirements	SAS 9.4 Intelligence Platform: Security Administration Guide
	Other compliance certifications	SAS Governance and Compliance Manager

Assumptions and prerequisites

To migrate SAS Grid to AWS, you must meet the assumptions and requirements discussed in this section. The migration of SAS software may require expert skills in SAS administration, system administration, and AWS administration. If you need help with determining the scope of migration for your SAS Grid environment, we recommend that you contact SAS professional services for the following assessments:

- Current SAS Grid workload assessment
- Security assessment
- SAS Grid migration assessment
- SAS Grid migration advisory service

Migration requirements

- The physical topology of source and target systems must be equivalent, including host machines and their roles, with the expectation that RAM, CPU, and disk volume/throughput will compare similarly. Also, source and target operating systems must be in the same family. For SAS installation prerequisites, see [SAS system requirements](#) on the SAS website.
- Source and target systems must be SAS 9.2 or later.
- Data, files, and other content that isn't migrated automatically must be migrated manually.
- This workload migration doesn't include original data providers. Rehosting original data on AWS, especially in a different data provider technology, requires additional effort.
- For SAS Bring Your Own License (BYOL) migration, you must establish and maintain the AWS environment.

Knowledge requirements

A solid understanding of the SAS system and the components of SAS infrastructure is required to optimize your SAS Grid environment on AWS. Considerations such as storage service, server instance types, networking performance, high availability, and disaster recovery affect the architecture design of your SAS environment on AWS.

Additional SAS considerations

- SAS infrastructure sizing and architecture must be created based on:
 - Instance types
 - Ephemeral, persistent, and shared storage types
 - A shared file system for SAS Grid Manager
 - Placement of SAS Permanent Data File Space (SASDATA) and temporary file spaces: SAS Working Data File Space (SASWORK) and SAS Utility Data File space (UTILLOC)
- SAS software licensing metrics are the same for SAS software cloud and on-premises deployments.
- Cloud administration, security, and monitoring are the responsibility of users, unless the environment has been contracted by SAS as part of a remotely managed environment.
- SAS software can be scaled, but you must be careful to comply with licensing agreements.

- In most cases, scaling a SAS infrastructure results in an outage of service during the process.
- High availability, disaster recovery, and backup and restoration are as important in SAS software cloud deployments as they are in SAS software on-premises deployments.
- Local laws and privacy regulations might affect the data you store in the cloud. For example, certain geographies might restrict the storage and processing of data in a cloud location out of country or state.
- The cost of a cloud infrastructure is a core consideration.

Cost and licensing

At a high level, the cost of migrating a SAS workload to AWS assumes establishing a new environment on AWS. The cost includes accounting for personnel time and effort as well as provisioning computing resources and licensing software for the new environment.

SAS licenses

Selecting the appropriate size instance type for the SAS license is an important consideration. In general, SAS licensing is based on the number of physical CPUs on the host machine(s). In comparison, AWS typically specifies the number of virtual CPUs (vCPUs, or hyperthreading) for its instance types. As a rule, for the target destination, select an Amazon Elastic Compute Cloud (Amazon EC2) instance that has **twice the number of vCPUs** as compared to the physical CPUs in the source environment. See [Optimizing CPU options](#) in the Amazon EC2 documentation for more information about optimizing vCPU-based licensing costs.

Another consideration is the cost of running two environments (source and target) at the same time, while getting the value of only one environment. SAS offers two policies to help reduce the cost of migration efforts like this, depending on the exact circumstances involved: policy for upgrading SAS software, and policy for testing new hardware.

Policy	Description	Duration
Upgrading SAS software	<p>SAS allows all customers who are upgrading to a new production release to run two parallel copies of the licensed software (the old release and the new release) for 180 days (6 months) at no additional charge.</p> <p>For metered offerings, customers can run two parallel copies of the licensed software, but both will be metered for billing purposes.</p>	180 days
Testing new hardware	<p>The following scenarios are covered by SAS evaluation paperwork and are offered free of charge for a limited test period. The standard test period for each scenario is 30 days.</p> <ul style="list-style-type: none">• Testing new hardware before upgrading production machines and production systems to the new hardware configuration; performing parallel processing.• Testing new operating systems before converting	30 days

Policy	Description	Duration
	from one operating system to another in a production environment; performing parallel processing.	

Note

Test machines are defined as hardware the customer is responsible for that is used to run SAS software in an evaluation environment to verify quality and accuracy of data output.

Contact your SAS account representative for more information specific to your site and license.

Cost estimates for AWS services

[AWS Pricing Calculator](#) is an online tool that helps you estimate the monthly cost of AWS services for your use case, based on your expected usage. The Pricing Calculator is continuously updated with the latest pricing for all AWS services in all Regions. The calculator includes support for most AWS services. You can include additional costs such as data ingress/egress charges, data storage charges, and retrieval fees. You can also select Amazon EC2 with various pricing models such as On-Demand, Dedicated, and Reserved Instances. We recommend that you use Pricing Calculator to estimate your monthly AWS bill more efficiently.

Note

Evaluate the ratio of vCPU to I/O bandwidth and vCPU to ephemeral storage per AWS instance type. Choose the instance that provides the best ratio of those two factors for your SAS Grid processing computing needs.

Workload sizing:

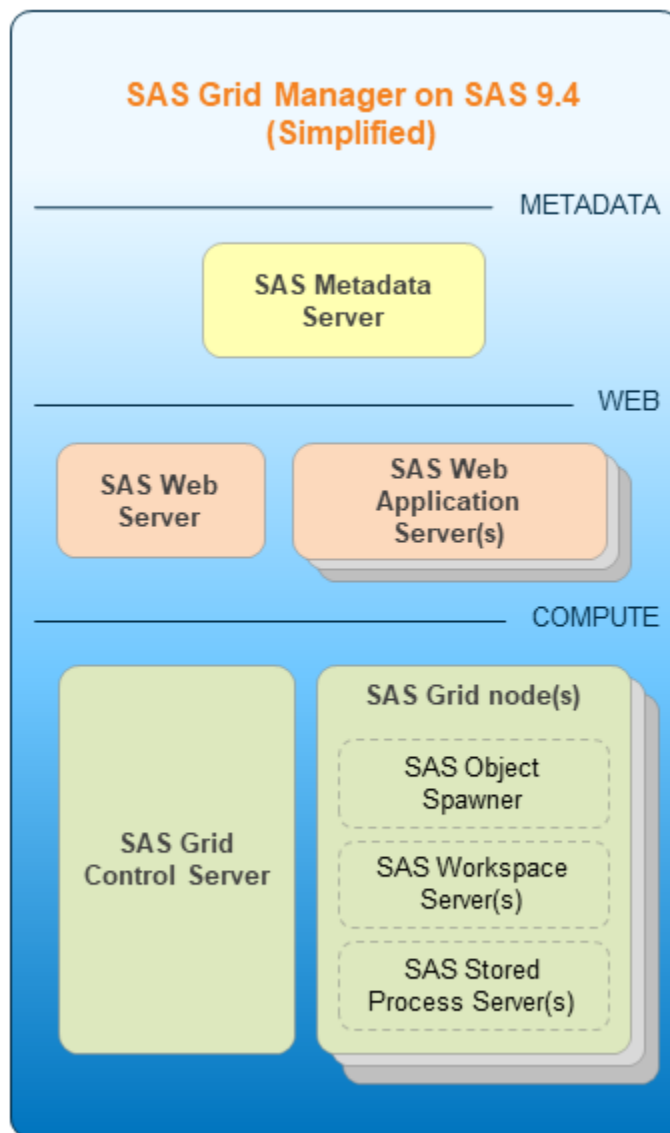
- **Small** – SAS Grid without high availability in a single Availability Zone
2x SAS Grid (i3en.6xlarge instance), 1x Metadata Server, 1x Web Server
- **Medium** – SAS Grid with high availability in a single Availability Zone
4x SAS Grid (i3en.12xlarge instance), 3x Metadata Server, 2x Web Server
- **Large** – SAS Grid with high availability in a single Availability Zone
8x SAS Grid (i3en.12xlarge instance), 3x Metadata Server, 2x Web Server

Architecture

Building the appropriate SAS Grid architecture is a critical step to ensure that migration and performance will meet your users' expectations. For your migration effort to meet this guide's assumptions, the target environment for SAS Grid Manager on AWS must be an equivalent match in terms of operating system and software versions as well as major machine specifications.

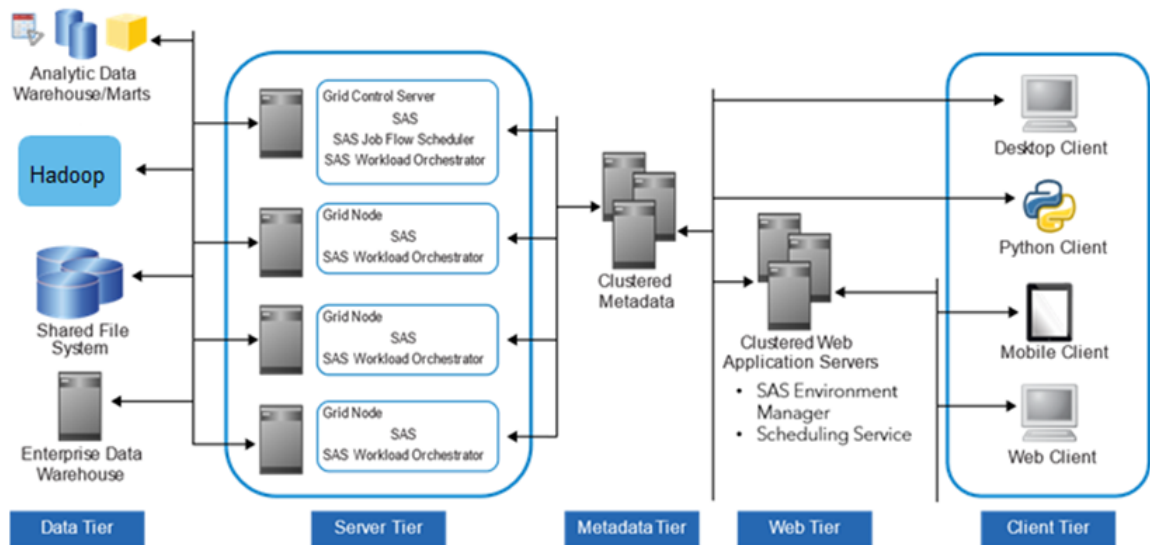
SAS infrastructure

The following diagram shows the infrastructure components of SAS Grid Manager. The illustration is simplified to highlight major components that provide end-user functionality or that must be considered when planning resource allocations for processing, memory, network, and I/O.



- **SAS Metadata Server** is the central hub of SAS Grid that client, server, and intermediate software components rely on. It provides information regarding software processes, manages user authentication and authorization to resources, and maintains user content.
- **SAS Web Server** hosts static collateral and also acts as a reverse proxy, providing a single point of contact to the web apps in their Java Virtual Machines (JVMs).
- **SAS Web Application Servers** host the various web apps for end-user access and operation, including SAS Studio, SAS Environment Manager, and others.
- SAS offers compute server processes that are specialized for their respective clients:
 - **SAS Object Spawner** initiates new SAS Integrated Object Model (IOM) processes.
 - **SAS Workspace Server** provides each user with a dedicated analytics environment for clients like SAS Enterprise Guide and SAS Studio.
 - **SAS Stored Process Server** acts a persistent analytics engine for predefined tasks (stored processes).
 - **SAS Grid Control Server** distributes jobs to one or more compute nodes on the grid. A grid control server can also do work allocated to the grid.
 - **SAS Grid nodes** run a portion of the work allocated to the grid.

The following architecture diagram shows how the tiers or infrastructure components interact.



Note

The five tiers represent categories of software that perform similar types of computing tasks and require similar types of resources. The tiers **do not necessarily** represent separate computers or groups of computers. For more information about each tier, use the links to the SAS documentation in the following list.

- **Data tier** – Stores your enterprise data. You can use all your existing data assets, including data stored in third-party database management systems, SAS tables, enterprise resource planning (ERP) system tables, and AWS-specific storage services such as Amazon FSx for Lustre or Amazon S3.
- **Server tier** – Performs SAS processing on your enterprise data. Several types of SAS servers are available to handle different workload types and processing intensities. The software distributes processing loads among server resources so that multiple client requests for information can be met without delay.

- **Metadata tier** – Client, server, and intermediate software components rely on SAS Metadata Server, which is the central hub of SAS Grid. It provides information regarding software processes, manages user authentication and authorization to access resources, and maintains user content.
- **Web tier** – Enables users to access intelligence data and functionality by using a web browser. This tier provides web-based interfaces for report creation and information distribution, and passes analysis and processing requests to the SAS servers.
- **Client tier** – Provides users with desktop access to intelligence data and functionality through easy-to-use interfaces. For most information consumers, reporting and analysis tasks can be performed with just a web browser. For more advanced design and analysis tasks, SAS client software is installed on users' desktops. Some support for mobile devices is also provided.

Considerations for SAS on AWS

SAS background

- SAS Permanent Data File Space (SASDATA)
 - Stores persistent data for SAS exploitation and resulting SAS output files
 - Extensively read from, but less extensively written back out
- SAS Working Data File Space (SASWORK)
 - Scratch working space for SAS jobs
 - Used to perform the working storage activity of single-threaded SAS procedures
- SAS Utility Data File Space (UTILLOC)
 - Same type of space as SASWORK for multi-threaded SAS procedures
 - By default, placed as a subdirectory under SASWORK
- Storage does not persist through a reboot or restart of internal solid state (SSD) devices that are striped together in a RAID 0 configuration. We recommend that you use instances with non-volatile memory express (NVMe) ephemeral devices with high bandwidth, low latency, and sequential I/O. These instances are ideal for temporary SAS data (SASWORK and UTILLOC).

SAS shared file system (required for SAS Grid)

- AWS sets up Lustre file systems with the **rw**, **seclabel**, and **lazystatfs** mount options. These are not the recommended mount options for SAS Grid, so you will have to unmount these FSx for Lustre file systems and remount them with the **flock** parameter.
- You cannot expand the size of your Lustre file system. To resize, create a larger Lustre file system and copy data from the old system to the new one.
- For FSx for Lustre persistent file systems, data is replicated within a single Availability Zone for increased durability. It is not replicated across AWS Availability Zones.
- We recommend that you use the Amazon S3 storage option for use with SAS Grid and FSx for Lustre. For more information, see [Using Data Repositories with Amazon FSx for Lustre](#) in the AWS documentation.
- See the [AWS Region table](#) for availability of services across AWS Regions and Availability Zones. Also review [Amazon S3 same-Region replication \(SRR\)](#) or [cross-Region replication \(CRR\)](#) for implications of data replication needs for high availability.

SAS Grid server tier instance types

SAS Grid servers need fast CPUs for processing data. We recommend:

- A minimum of 8 GB of physical RAM per physical core, and robust I/O throughput (especially for SASWORK and SAS UTILLOC).
- [I3 instances](#) – Amazon EC2 I3 instances are storage-optimized for high transaction, low latency workloads. These instances include NVMe SSD-based instances that are storage-optimized for high random I/O performance, high sequential read throughput, and high IOPS. Because of the high internal I/O bandwidth from striped NVMe SSD drives for SASWORK and SAS UTILLOC, you should configure your environment to explicitly use the NVMe-based SSD local drives instead of Amazon EBS volumes.
- [I3en instances](#) – This family provides NVMe SSD instances that are storage-optimized on Amazon EC2 with enhanced networking via ENA to achieve up to 100 Gbps of network bandwidth.
- [M5n instances](#) – The M5 family provides a balance of compute, memory and networking. M5n instances are ideal for applications that require improved network throughput and packet rate performance.
- SAS workloads can be characterized as predominately large, sequential I/O requests with high volumes of data. We recommend that you predetermine your SAS usage patterns. This will guide the optimal architecture and setup of the individual underlying file systems and their respective physical I/O provisioning.
 - Query, reporting, and light statistical jobs usually perform well with an I/O rate of 100 MiB per second per physical CPU core.
 - Advanced analytics and heavy statistical jobs might require up to 150 MiB per second per physical CPU core.
 - Overall, we recommend a minimum I/O throughput rate of 100-125 MiB per second per physical CPU core.

SAS Grid middle tier and metadata server tier instance types

These servers do not require computing-intensive resources or robust I/O throughput. They do require access to more memory than the SAS computing tiers. We recommend:

- A minimum of 24 GB of physical RAM or 8 GB of physical RAM per physical core, whichever is larger.
- [R5 or R5d instances](#) – These instances are suitable for memory-intensive applications such as in-memory caches, mid-size in-memory databases, and real-time big data analytics.

High availability and disaster recovery for SAS Grid

Disaster recovery planning is important for any critical business system, including production systems that are running the SAS Intelligence Platform and SAS solutions.

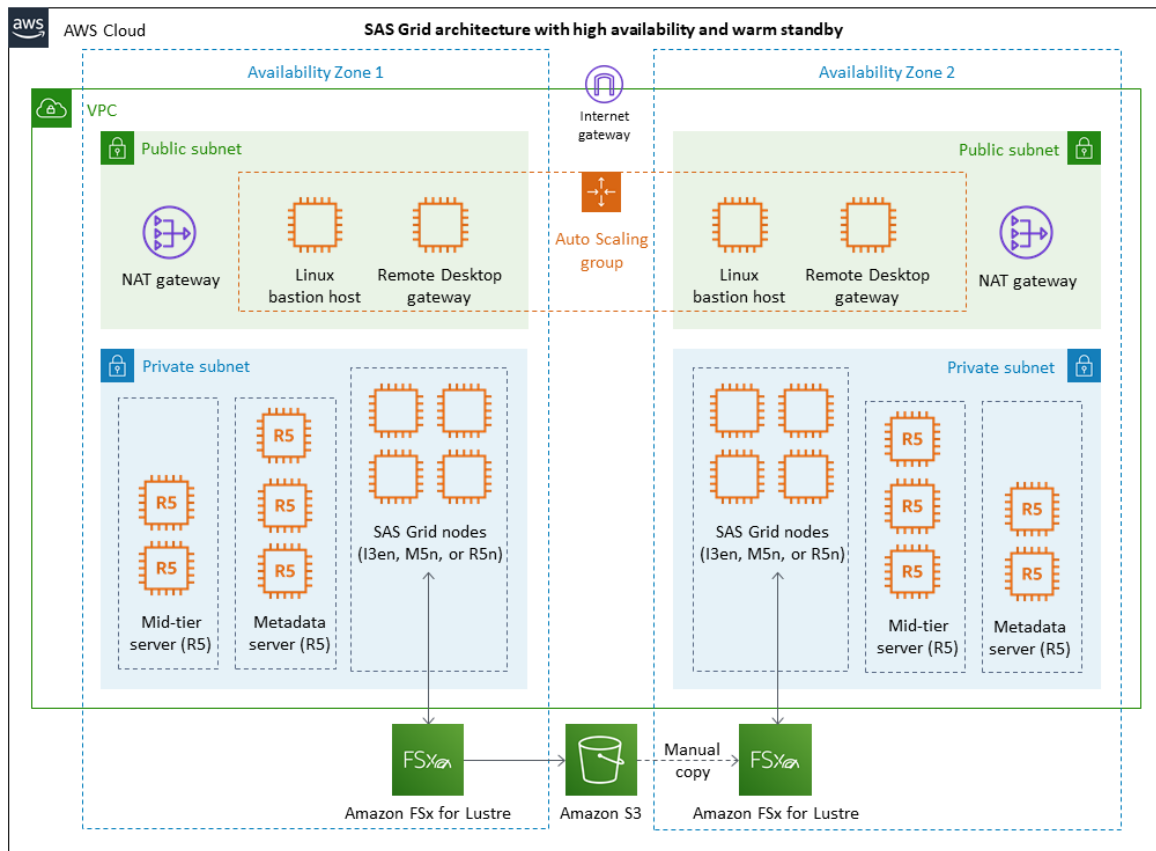
Disaster recovery is not the same as high availability. Although both concepts are related to business continuity, high availability is about providing uninterrupted continuity of operations. In contrast, disaster recovery involves some amount of downtime, typically measured in hours or days.

Target architecture

Although you can choose the appropriate instance types for your specific workload needs, for SAS Grid Manager on SAS 9.4, SAS recommends [Amazon EC2 I3en instances](#). We also recommend using [Amazon VPC](#), which provides increased isolation control, customization, and security.

AWS Prescriptive Guidance Migrating
SAS Grid to the AWS Cloud
Target architecture

The following diagram shows SAS Grid on AWS with data, metadata, middle tier, and server tiers. This high-availability architecture is deployed on two Availability Zones for an active-active disaster recovery failover strategy.



This architecture includes the following components:

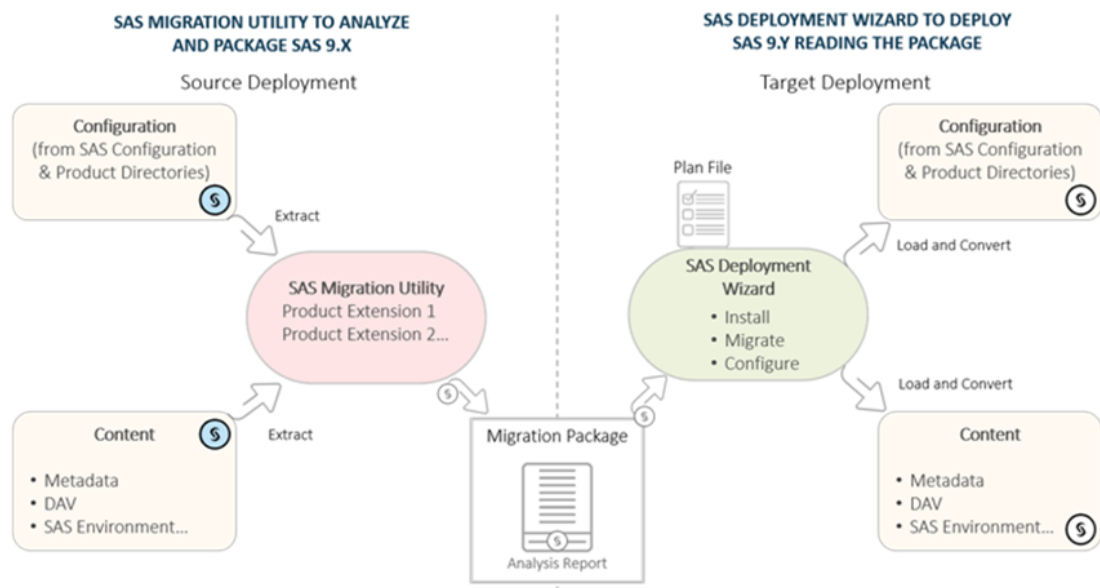
- **Virtual private cloud (VPC)** – A virtual network dedicated to your AWS account. This is logically isolated from other virtual networks in the AWS Cloud. You can also create a hardware virtual private network (VPN) connection between your corporate data center and your VPC, and use the AWS Cloud as an extension of your corporate data center. The VPC is configured with two Availability Zones, public subnets, and private subnets to provide the network infrastructure for SAS Grid on AWS.
- **Internet gateway** – This gateway is attached to your VPC. By default, it comes with a security group that allows **no inbound** traffic and **all outbound** traffic to the internet.
- **NAT gateway** – The network address translation (NAT) gateways enable instances in private subnets to connect to the internet.
- **Linux bastion host** – Provides secure access to Linux instances located in the private and public subnets of your VPC.
- **Remote Desktop gateway** – Provides remote administration. This gateway uses remote desktop protocol (RDP) over HTTPS to establish a secure, encrypted connection between remote users on the internet and Windows-based EC2 instances.
- **Amazon EC2 Auto Scaling** – Ensures that the number of bastion hosts and Remote Desktop gateway instances always matches the capacity you specify during launch.
- **FSx for Lustre** – Integrates with Amazon S3 and makes it easy to process cloud datasets using the Lustre high-performance file system.

- [Amazon S3](#) – Enables you to store and retrieve any amount of data at any time, from anywhere on the web.

Automation and tooling

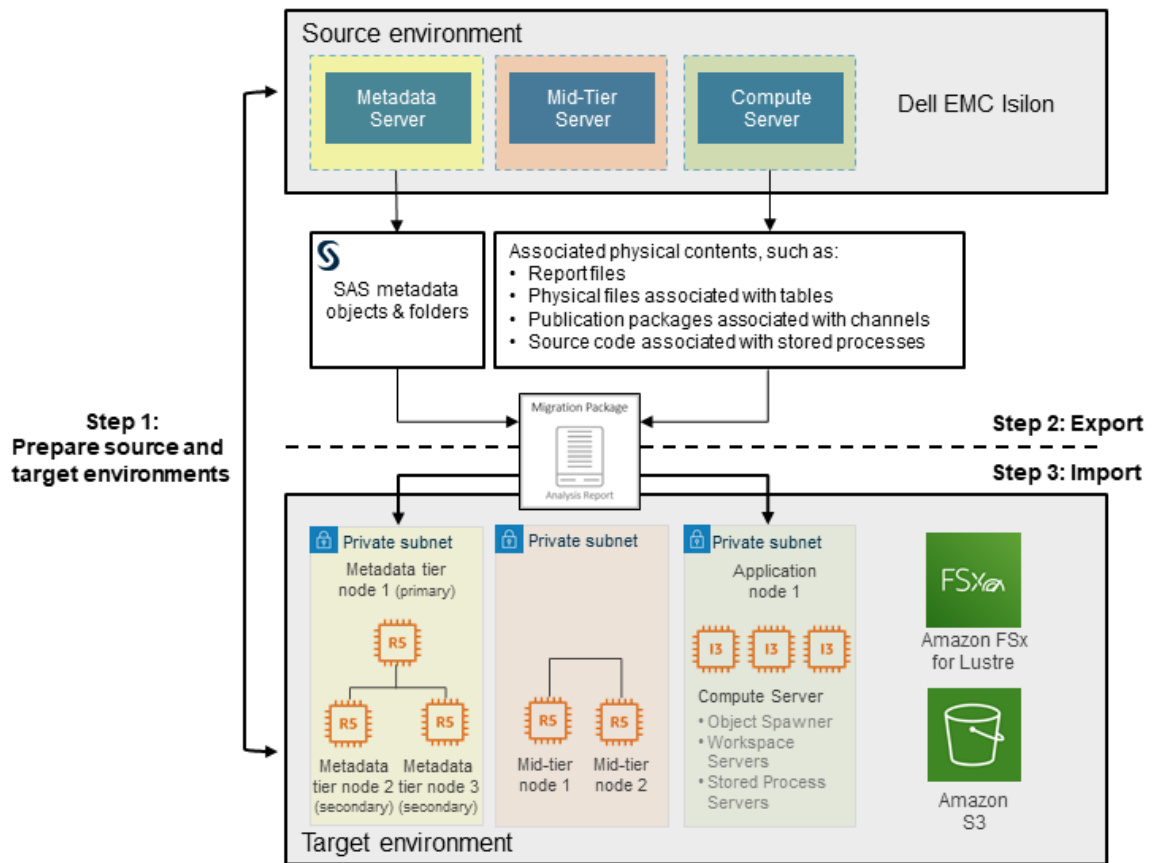
Matching source and target topologies simplifies many variables that can affect migration efforts. In addition, matching topologies is a requirement for using the SAS Migration Utility. This tool assumes that every host machine, directory, and network component in the source environment will map one-to-one to its equivalent in the target environment.

You use the SAS Migration Utility to analyze and package your source environment. As the following diagram shows, the resulting package file is copied to the target system you have provisioned on AWS. It is then processed by the SAS Deployment Wizard as part of the initial SAS Enterprise BI Server software deployment on AWS.



The SAS Migration Utility is useful for migrating SAS metadata content and certain associated files stored in the configuration directory. The bulk of the physical files (SAS data sets, programs, external files, and so on) aren't part of the SAS Migration Utility process and have to be copied over to the target environment separately. To copy these physical files over, we recommend that you explore the use of [AWS DataSync](#).

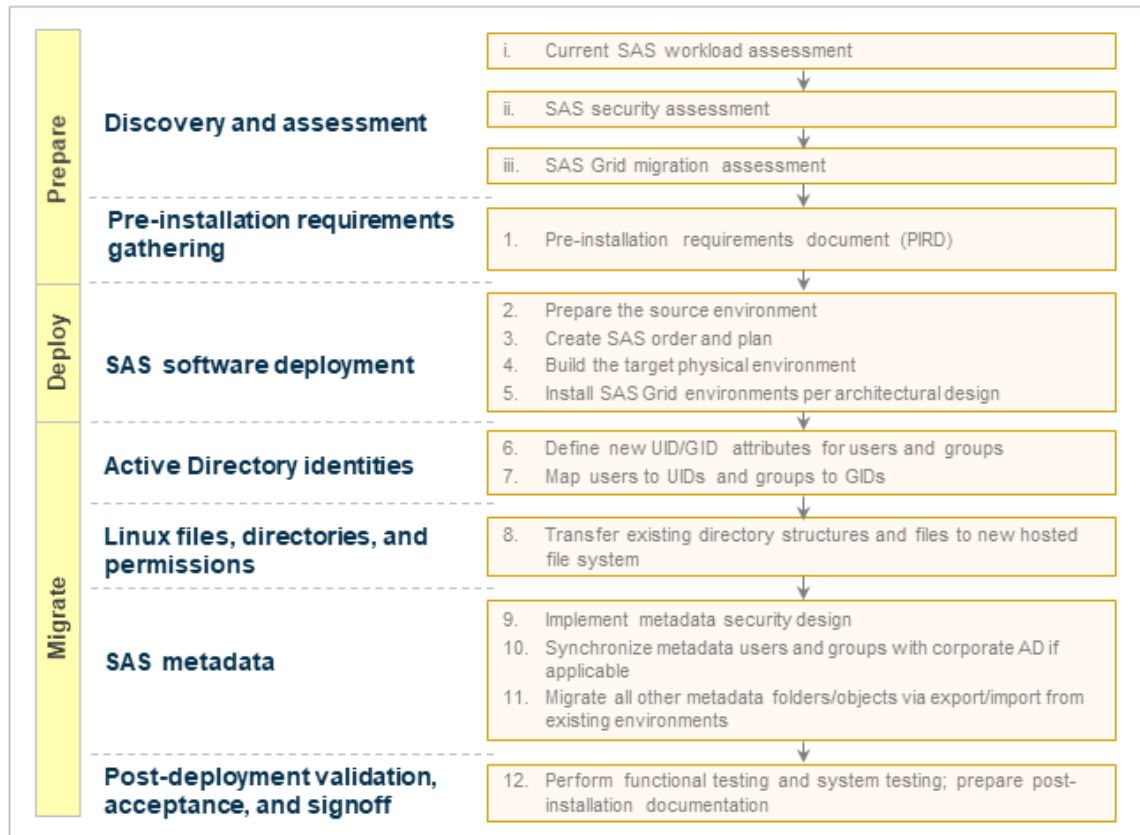
After the majority of SAS content has been migrated to the new target system on AWS, you perform validation testing. Before you cut over to your production environment on AWS, you can perform one final promotion of any content that has been added or changed since the last migration.



For additional details on the SAS Migration Utility, see the [SAS 9.4 Intelligence Platform Migration Guide](#) on the SAS website.

High-level migration steps

The following table describes the steps for migrating SAS Grid to AWS, organized by epics and stories. These classifications are used in agile methodologies to help you scope and manage your work. Each epic is described in detail in subsequent sections.



Roles and responsibilities

The SAS Grid migration process involves the following roles and expertise.

Role	Responsibilities
SAS SME or consultant	<ul style="list-style-type: none">• Provide SAS installation, technical architecture, security, and migration expertise.
AWS architect	<ul style="list-style-type: none">• Advise and provision AWS services and infrastructure resources.• Establish AWS Identity and Access Management (IAM) roles for users and groups.• Provide security administration for all provisioned AWS services.

Prepare – Discovery and assessment

Prepare current SAS Grid workload assessment

- Goal: Capture the infrastructure requirements for SAS Grid on AWS.
- Tasks:
 - Gather RHEL performance information by running the [gather_info.sh script](#) on every system in the SAS infrastructure.
 - Test throughput for your SAS file systems by running the [rhel_iotest.sh script](#).
 - Use the [nmon command](#) to monitor the performance of SAS on a Linux system.
 - Run [nmon](#) for several days for an average to above-average workload.
 - Collect data every 5 minutes for 24 hours.
 - Return all logs and files to SAS for analysis.
 - Analyze current utilization patterns and present infrastructure estimate in AWS.
- Skills/roles: SAS consultant, SAS technical support, SAS Performance Lab

Prepare SAS security assessment

- Goal: Gather requirements to create SAS security design documentation.
- Tasks:
 - Gather requirements:
 - Define functional security requirements for the SAS environment.
 - Determine whether compliance to an external standard is required.
 - Define user personas such as administrators, analysts, and report viewers.
 - Define organizational hierarchy, and determine how content will be organized and secured.
 - Document SAS security design, including:
 - Metadata authorization model, which defines groups, roles, folders, server metadata, and access control templates.
 - File system security design, which identifies high-level, key SAS directories; the location, ownership, and permissions of those directories; and special files in those directories.
 - SAS security design specification document.
- Skills/roles: SAS consultant

Prepare SAS Grid migration assessment

- Goal: Review high-level SAS migration requirements, and use the SAS Migration Utility to analyze your SAS content and validate that it can be migrated successfully.
- Tasks:
 - Review [high-level SAS migration requirements](#).
 - Inventory the source SAS deployment:
 - [Complete a migration utility checklist](#).
 - [Check backups in SAS](#).

- [Prepare your metadata repositories.](#)
- [Review SAS Migration Utility requirements.](#)
- [Download the SAS Migration Utility.](#)
- [Develop a common SAS Migration Utility properties file.](#)
- [Review product-specific SAS Migration Utility properties.](#)
- [Create a migration analysis report.](#)
- [View and analyze your migration analysis report and migration log.](#)
- Skills/roles: SAS consultant

Set up a new AWS environment (for new AWS users only)

- Goal: If you're a SAS customer who is new to AWS, you can use [AWS Control Tower](#) to set up a base environment. To migrate SAS Grid to AWS, you must add a VPC to your existing environment at a minimum.
- Tasks:
 - Use AWS Control Tower to set up a well-architected, multi-account environment that is based on security, compliance, and other AWS best practices. You could work within a single account. However, as your environment grows, it will be easier to manage it through a multi-account infrastructure.
 - When the new AWS environment is set up, configure authentication:
 - For Windows servers, use [AWS Directory Service](#) or [Active Directory Connector](#) to connect to your existing on-premises Microsoft Active Directory. For added security, use multi-factor authentication (MFA). For more information about using AD Connector with MFA, see [Enable MFA for AD Connector](#) in the AWS Directory Service documentation.
 - For Linux instances, you can use AWS Directory Service with certain EC2 Linux instances. For details, see [Manually Join a Linux Instance](#) in the AWS Directory Service documentation.
 - If on-premises access required, see [Access to an on-premises network](#) in the AWS Client VPN documentation.
- Skills/roles: AWS architect

Deploy – SAS software on AWS

- Goal: Install and configure SAS software on AWS. For more information, see [Installing SAS 9.4 and Migrating Your SAS Content](#) in the SAS documentation.
- Tasks:
 - Prepare the source environment.
 - Create an SAS order and plan.
 - Build the target physical environment.
 - Install the SAS Grid environments in accordance with your architectural design.
- Skills/roles: SAS consultant

Migrate – SAS content to AWS

Migrate Active Directory identities

- Goal: Define user IDs (UIDs) and group IDs (GIDs), and configure mapping.
- Tasks:
 - Define new UID/GID attributes for users and groups.
 - Map users to UIDs and security groups to GIDs.
- Skills/roles: SAS consultant

Migrate Linux files, directories, and permissions

- Goal: Transfer Linux files and directories, while maintaining the directory structure on AWS.
- Tasks:
 - Transfer your existing directory structures and files to the file system on AWS. (For example, you can use [Amazon FSx for Lustre](#).)
 - Make sure that the directory structure and file paths are consistent in the source and target environments to minimize code changes.
- Skills/roles: SAS consultant

Migrate SAS metadata

- Goal: Implement metadata security design and synchronize SAS metadata content.
- Tasks:
 - Implement SAS metadata security design on AWS.
 - Synchronize users and groups from your existing corporate Active Directory, where appropriate, based on the security design.
 - Migrate all other SAS metadata folders and objects via export/import from your existing environment.
- Skills/roles: SAS consultant

Perform post-migration validation and acceptance testing

- Goal: Perform functional and system testing, sign off on the migration, and create post-migration reports. For details, see [Performing Post-Migration Tasks](#) in the SAS documentation.
- Tasks:
 - Perform functional testing of the SAS application on AWS.
 - Perform system testing of SAS applications on AWS.
 - Prepare post-installation documentation.

AWS Prescriptive Guidance Migrating
SAS Grid to the AWS Cloud
Perform post-migration validation and acceptance testing

- Skills/roles: SAS consultant

Migrating data associated with SAS

We recommend that you move the data associated with your SAS applications to AWS. This migration has several benefits:

- Gaining access to cloud-based data lakes and data warehouses
- Increased agility, performance, security, and reliability
- Lower costs

AWS offers a wide variety of services and tools to help you migrate your data sets, including SAS files, databases, machine images, block volumes, and even tape backups. The following table provides a list of services that you can use.

AWS service	Description	Role/skills required
AWS DataSync	Copies or replicates file system data to Amazon S3 or Amazon Elastic File System (Amazon EFS).	AWS architect
CloudEndure Migration	Migrates running machine images with their data to Amazon EC2.	AWS architect
Amazon S3 Transfer Acceleration	Enables fast and secure transfers of data to Amazon S3 over long geographic distances.	AWS architect
AWS DMS	Migrates databases to AWS quickly and securely, with minimal downtime.	AWS architect
AWS Snow Family	Physically transports petabytes of data in batches to AWS.	AWS architect

Onboarding new customers

When migration is complete, you'll want to onboard new and existing users to SAS Grid on AWS. This might also involve training for SAS end-users or administrators. We recommend the following tasks:

- SAS Grid administration onboarding and training (if appropriate)
- SAS end-user onboarding and training (if appropriate)
- SAS customer learning portal onboarding (if appropriate)

To explore onboarding options and resources, see the [SAS Self-Paced Digital Learning Portal](#).

Additional resources

SAS papers

- [Important Performance Considerations When Moving SAS to a Public Cloud](#)
- [Best Practices for Upgrading from SAS 9.1.3 to SAS 9.4](#)
- [Getting the Latest and Greatest from SAS 9.4: Best Practices for Upgrades and Migrations](#)

SAS documentation

- [Grid Computing in SAS 9.4, Fifth Edition](#)
- [SAS 9.4 Intelligence Platform: Overview, Second Edition](#)
- [SAS 9.4 Intelligence Platform: Migration Guide, Second Edition](#)
- [SAS 9.4 Foundation for Linux for x64](#)
- [SAS Studio 3.71](#)
- [SAS Enterprise Guide 8.2](#)

AWS documentation and resources

- [AWS Region Table](#)
- [Cloud Data Migration](#)
- [AWS DMS](#)
- [Amazon EC2](#)
- [Amazon FSx for Lustre](#)
- [IAM](#)
- [Amazon S3](#)
- [Amazon VPC](#)

AWS Prescriptive Guidance glossary

[AI and ML terms \(p. 25\)](#) | [Migration terms \(p. 26\)](#) | [Modernization terms \(p. 29\)](#)

AI and ML terms

The following are commonly used terms in artificial intelligence (AI) and machine learning (ML)-related strategies, guides, and patterns provided by AWS Prescriptive Guidance. To suggest entries, please use the **Provide feedback** link at the end of the glossary.

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?"
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.
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.
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.
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.
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 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.

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?"
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).
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.
target variable	The value that you are trying to predict in supervised ML. This is also referred to as an <i>outcome variable</i> . For example, in a manufacturing setting the target variable could be a product defect.
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.
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: <i>Epistemic uncertainty</i> is caused by limited, incomplete data, whereas <i>aleatoric uncertainty</i> is caused by the noise and randomness inherent in the data. For more information, see the Quantifying uncertainty in deep learning systems guide.

Migration terms

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

7 Rs	<p>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:</p> <ul style="list-style-type: none">• 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. This migration scenario is specific to VMware Cloud on AWS, which supports virtual machine (VM) compatibility and workload portability between your on-premises environment and AWS. You can use the VMware Cloud Foundation technologies from your on-premises data centers when you migrate your infrastructure to VMware Cloud on AWS. Example: Relocate the hypervisor hosting your Oracle database to VMware Cloud on 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.

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 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](#).

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 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](#).

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.

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.

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-

	<p>scale transformations. For more information, see the CCoE posts on the AWS Cloud Enterprise Strategy Blog.</p>
cloud stages of adoption	<p>The four phases that organizations typically go through when they migrate to the AWS Cloud:</p> <ul style="list-style-type: none">• 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 <p>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.</p>
configuration management database (CMDB)	<p>A database that contains information about a company’s hardware and software products, configurations, and inter-dependencies. You typically use data from a CMDB in the portfolio discovery and analysis stage of migration.</p>
epic	<p>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 execution guide.</p>
heterogeneous database migration	<p>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.</p>
homogeneous database migration	<p>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.</p>
IT information library (ITIL)	<p>A set of best practices for delivering IT services and aligning these services with business requirements. ITIL provides the foundation for ITSM.</p>
IT service management (ITSM)	<p>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.</p>
Migration Acceleration Program (MAP)	<p>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.</p>
Migration Portfolio Assessment (MPA)	<p>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</p>

	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 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 <i>migration factory</i> 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 CloudEndure Migration Factory guide in this content set.
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).
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 .
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 <i>people acceleration</i> , because of the speed of change required in cloud adoption projects. For more information, see the OCM guide .
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.
responsible, accountable, consulted, informed (RACI) matrix	A matrix that defines and assigns roles and responsibilities in a project. For example, you can create a RACI to define security control ownership or to identify roles and responsibilities for specific tasks in a migration project.
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.
service-level agreement (SLA)	An agreement that clarifies what an IT team promises to deliver to their customers, such as service uptime and performance.

Modernization terms

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

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.
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 .
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 .
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 .
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 .
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.

For more information, see the [Two-pizza team](#) section of the [Introduction to DevOps on AWS](#) whitepaper.

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](#).

update-history-change	update-history-description	update-history-date
Initial publication (p. 32)	—	July 17, 2020