



Creating an enterprise encryption strategy for data at rest

AWS Prescriptive Guidance



AWS Prescriptive Guidance: Creating an enterprise encryption strategy for data at rest

Copyright © 2024 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

Introduction	1
Intended audience	1
Targeted business outcomes	2
Limitations	2
About data encryption	4
About encryption keys	4
About encryption algorithms	4
About envelope encryption	4
Encryption strategy phases	6
Policy	6
Standards	7
Cost and performance	8
Key access control	8
Encryption types	9
Encryption key specifications	9
Key storage location	9
Framework	10
Data classification	10
Environment classification	11
Change events and processes	11
Implementation	12
Cost, convenience, and control	13
Performance and encryption types	14
Key storage location	14
Access control	15
Auditing and logging	16
FAQ	17
When do I need symmetric encryption?	17
When do I need asymmetric encryption?	17
When do I need envelope encryption?	17
When do I need to use an HSM?	17
Why should I centrally manage encryption keys?	18
Do I need to use a purpose-built encryption infrastructure?	18
How can AWS KMS help?	18

Resources	20
AWS service documentation	20
AWS marketing	20
AWS Well-Architected Framework	20
Hashing and tokenization	20
Videos	21
Document history	22
Glossary	23
#	23
A	24
B	27
C	29
D	32
E	36
F	38
G	39
H	40
I	41
L	43
M	44
O	48
P	51
Q	53
R	54
S	56
T	60
U	61
V	62
W	62
Z	63

Creating an enterprise encryption strategy for data at rest

Venki Srivatsav, Andrea Di Fabio, and Vikramaditya Bhatnagar, Amazon Web Services (AWS)

September 2022 ([document history](#))

Many enterprises are concerned about the cybersecurity threat of a data breach. When a data breach occurs, an unauthorized person gains access to your network and steals enterprise data. Firewalls and anti-malware services can help protect against this threat. Another protection that you can implement is data encryption. In the About data encryption section of this guide, you can learn more about how data encryption works and the types available.

When you're discussing encryption, generally speaking, there are two types of data. *Data in transit* is data that is actively moving through your network, such as between network resources. *Data at rest* is data that is stationary and dormant, such as data that is in storage. This strategy focuses on data at rest. For more information about encrypting data in transit, see [Protecting data in transit](#) (AWS Well-Architected Framework).

An *encryption strategy* consists of four parts that you develop in sequential phases. The *encryption policy* is determined by senior management and outlines the regulatory, compliance, and business requirements for encryption. The *encryption standards* help those who implement the policy to understand it and comply with it. Standards can be technological or procedural. The *framework* is the standard operating procedures, structures, and guardrails that support implementation of the standards. Finally, the *architecture* is the technical implementation of your encryption standards, such as the environment, services, and tools you use. The objective of this document is to help you create an encryption strategy that suits your business, security, and compliance needs. It includes recommendations for how to review and implement security standards for data at rest so that you can meet your compliance and business needs in a holistic manner.

This strategy uses AWS Key Management Service (AWS KMS) to help you create and manage cryptographic keys that help protect your data. AWS KMS integrates with many AWS services to encrypt all your data at rest. Even if you choose a different encryption service, you can still adopt the recommendations and phases in this guide.

Intended audience

The strategy is designed to address the following audiences:

- Executive officers who formulate policies for their enterprise, such as CEOs, chief technology officers (CTOs), chief information officers (CIOs), and chief information security officers (CISOs)
- Technology officers who are responsible for setting up technical standards, such as technical vice presidents and directors
- Compliance and governance officers who are in charge of monitoring adherence to compliance policies, including statutory and voluntary compliance regimes

Targeted business outcomes

- **Data-at-rest encryption policy** – Decision and policy makers can create an encryption policy and understand the critical factors that affect the policy.
- **Data-at-rest encryption standards** – Technical leaders can develop encryption standards that are based on the encryption policy.
- **Framework for encryption** – Technical leaders and implementers can create a framework that acts as a bridge between those who determine the policy and those who create the standards. Framework, in this context, means identifying the appropriate process and workflow that helps you implement the standards within the confines of the policy. A framework is similar to a standard operating procedure or a change management process for changing policies or standards.
- **Technical architecture and implementation** – Hands-on implementers, such as developers and architects, are aware of the available architecture references that can help them implement the encryption strategy.

Limitations

This document is intended to help you formulate a custom encryption strategy that best suits your enterprise's needs. It isn't an encryption strategy itself, and it isn't a compliance checklist. The following topics aren't included in this document:

- Encrypting data in transit
- Tokenization
- Hashing
- Compliance and data governance
- Budgeting for your encryption program

For more information about some of these topics, see the [Resources](#) section.

About data encryption

This section contains a high-level overview of encryption concepts and terminology. For more detailed information, see [Cryptography concepts](#) (AWS cryptographic services and tools guide). Data encryption helps you enforce data confidentiality. By implementing encryption and access controls, you can help protect the data in your enterprise.

About encryption keys

Encryption services use an encryption key to encrypt data. An *encryption key* is a cryptographic string of randomized bits that is generated by an encryption algorithm. Keys can vary in length, and each key is designed to be unpredictable and unique. The strength of the encryption typically depends on two factors: the length of the key and the algorithm used. In general, longer keys provide stronger encryption.

About encryption algorithms

There are two types of algorithms for generating encryption keys, symmetric and asymmetric.

Symmetric encryption uses the same key to encrypt and decrypt the data. This type of encryption is typically faster and is, therefore, efficient for large amounts of data. This type of encryption is widely used and generally accepted to be secure. Because a single key is used for both encryption and decryption, the best practice is to change the key frequently to prevent an unauthorized person from obtaining it. For more information about when symmetric encryption is recommended, see [When do I need symmetric encryption?](#) in the *FAQ* section.

Asymmetric encryption uses a pair of keys, a public key for encryption and a private key for decryption. You can share the public key because it isn't used for decryption, but access to the private key should be highly restricted. Asymmetric encryption is generally regarded to be more secure than symmetric encryption, but it is slower because it uses longer key lengths and requires more complex encryption calculations. For more information about when asymmetric encryption is recommended, see [When do I need asymmetric encryption?](#) in the *FAQ* section.

About envelope encryption

When you encrypt your data, it is protected only as long as your encryption key remains secret. The key used to encrypt the data is known as a *data key*. *Envelope encryption* is the practice of

encrypting your data key with another encryption key, called a *key-encryption key*. You can even encrypt that key with another encryption key, and so on. Eventually, one key must remain in plaintext so you can decrypt the keys and your data. This top-level plaintext key encryption key is known as the *root key*.

Envelope encryption offers several benefits:

- **Convenience** – Because your data key is encrypted, you can store it with the encrypted data.
- **Efficiency** – Encryption operations can be time consuming, particularly when it's a large amount of data. Instead of re-encrypting raw data multiple times with different keys, you can re-encrypt only the data keys that protect the raw data. This allows you to provide two or more layers of encryption protection without re-encrypting the data.
- **Performance** – You can combine encryption algorithms. For example, you can use symmetric encryption for the raw data but use asymmetric encryption for the data key, which combines the strengths of both encryption algorithms.

For more information about envelope encryption, see [Envelope encryption](#) (AWS Key Management Service documentation). For more information about deciding if you need envelope encryption, see [When do I need envelope encryption?](#) in the *FAQ* section.

Phases of building an encryption strategy

Building an enterprise-level encryption strategy requires a multi-phased approach. Each phase defines a set of controls to help you achieve your desired, tangible results. This document guides you through these phases and asks you specific questions to help you customize your encryption strategy.

Building an encryption strategy for data at rest consists of the following sequential phases:

1. [Encryption policy](#) – Build a policy that defines the data-at-rest encryption objectives for your enterprise.
2. [Encryption standards](#) – Define the technical and procedural standards that help you realize your enterprise policy.
3. [Encryption framework](#) – Build the framework that helps all stakeholders understand, change, and implement your encryption standards.
4. [Implementation](#) – Deploy your encryption infrastructure.

Encryption policy

The purpose of an encryption policy is to establish, at a senior management level, the business and compliance expectations that the organization needs to meet. The policy serves as a starting point to define a suitable encryption strategy. The policy should be abstract enough to provide freedom and flexibility for implementation. At the same time, it must be specific enough to define the confines of an acceptable implementation that meets organizational objectives. In general, policies are technology-agnostic and very infrequently changed because they define the fundamental characteristics of your enterprise encryption strategy.

Typically, encryption policies contain, but are not limited to, the following:

- Any regulatory or compliance regimes that your enterprise must meet
- Any business commitments or expectations for data encryption
- The type of data that must be encrypted
- Criteria for when to use data-protection techniques other than encryption, such as hashing or tokenization

The highest management level of the organization, such as the CIO, CTO, and CISO, usually define and approve the encryption policy.

Consider the following when creating your encryption policy:

- Your line of business determines the compliance and regulatory regimes you need to adhere to. These regimes dictate the data encryption requirements. Executive-level decisions to expand the business into new regions or expand product offerings can affect which regulations apply for your data. For example, if a bank decides to offer credit cards to its customers, they probably need to be compliant with the [payment card industry Data Security Standard](#) (PCI-DSS), which requires data encryption.
- Your policy should specify what type of data needs to be encrypted. This varies based on compliance requirements and the data-handling objectives of your enterprise. For example, your policy might state that any data that the business captures or owns must be encrypted at rest.
- Your encryption policy must align with your internal data categorization standards. To formulate an effective encryption policy, determination of data categories at the metadata level is required. For example, your categories might include public, internal, confidential, secret, or customer data.
- Include criteria for how to determine which data should be encrypted and which data should be protected with another technique, such as tokenization or hashing. For example, your policy might state *Any personally identifiable information (PII) that goes to the audit, trace, or application logs must be tokenized.*

Encryption standards

Standards are derived from your policy. These are narrower in scope and help define the framework and architecture for implementation. For example, if your organization's policy is to encrypt your data at rest, then a standard would define what type of encryption is required and provide general direction about how to adhere to the policy.

Encryption standards commonly specify the following:

- The types of encryption that should be used
- Minimum specifications for encryption keys
- Who has access to encryption keys
- Where encryption keys should be stored

- Criteria for picking an appropriate key strength when choosing encryption or hashing techniques
- Key rotation frequency

Whereas you rarely need to update an encryption policy, encryption standards are subject to change. The cybersecurity industry constantly evolves to meet the ever-changing threat landscape. As such, your standards should change to adopt the latest technologies and best practices in order to provide the best possible protection for your enterprise data.

In an enterprise organization, vice presidents, directors, or data stewards typically define encryption standards, and a compliance officer typically reviews and approves them.

Consider the following categories of factors when defining and maintaining encryption standards in your organization:

- [Cost and performance considerations](#)
- [Key access control](#)
- [Encryption types](#)
- [Encryption key specifications](#)
- [Key storage location](#)

Cost and performance considerations

Consider the following operational factors when determining encryption standards for data at rest:

- The available hardware resources must be able to support your standards at scale.
- The cost of encryption varies based on the length of the key, the amount of data, and the time required to perform the encryption. For example, when compared to symmetric encryption, asymmetric encryption uses longer keys and takes more time.
- Consider the performance requirements of your enterprise applications. If your application requires low latency and high throughput, then you might want to use symmetric encryption.

Key access control

Identify access control policies for your encryption keys based on the principle of least privilege. *Least privilege* is the security best practice of granting users the minimum access they need to perform their job functions. In your standards, define an access control policy that:

- Identifies the roles that manage the key-encryption keys and data keys.
- Defines and maps key permissions to roles. For example, it defines who has key admin privileges and who has and key user privileges. Key admins can create or modify key-encryption keys, and key users can encrypt and decrypt data and generate data keys.

Encryption types

In your standards, define which encryption types and features are suitable for your organization:

- Document when to use symmetric and asymmetric encryption algorithms. For more information, see [When do I need symmetric encryption?](#) and [When do I need asymmetric encryption?](#) in the *FAQ* section.
- Decide whether you should use envelope encryption, and define the circumstances. For more information, see [When do I need envelope encryption?](#) in the *FAQ* section.
- Define criteria for when to use encryption alternatives, such as tokenization and hashing.

Encryption key specifications

Define required specifications for your encryption keys, such as key strength and algorithms. These specifications must comply with the regulatory and compliance regimes defined in the policy.

Consider defining the following specifications:

- Define the minimum key strength and algorithms for both symmetric and asymmetric encryption types. The factors of key strength include the length, randomness, and uniqueness.
- Define when you want to implement new versions of encryption algorithms. For example, your standards might state *Implement the latest version of the algorithm within 30 days of release* or *Always use one version older than the latest release*.
- Define the interval for rotating your encryption keys.

Key storage location

In your standards, consider the following when deciding where to store your encryption keys:

- Compliance and regulatory requirements might dictate where your encryption keys can be stored.

- Decide whether you want to store keys in a centralized location or with their corresponding data. For more information, see [Why should I centrally manage encryption keys?](#) in the *FAQ* section.
- If you choose centralized storage, decide whether to store keys in an enterprise-managed infrastructure, such as a hardware security module (HSM), or a managed service provider, such as AWS Key Management Service. For more information, see [When do I need to use a hardware security module \(HSM\)?](#) in the *FAQ* section.

Encryption framework

A *framework*, in this context, refers to a set of standard operating procedures that need to be followed when you modify the encryption standards or policy. The framework is the scaffolding that helps you implement the standards. It helps convert words into actions. The framework links the people who define standards with the people who implement them.

Frameworks typically include the following topics:

- [Data classification](#)
- [Environment classification](#)
- [Change events and processes](#)

Data classification

Data classification plays a vital role in creating an encryption strategy. *Data classification* is the process of assigning data to a category based on the sensitivity of the data. The following are common data classification categories, in increasing order of sensitivity: public, private, internal, confidential, and restricted.

Your encryption framework should include the following information about data classification:

- The data classification categories for your enterprise.
- The classification criteria used to classify data into its appropriate category. For example, a company's trade recipe could be classified as *restricted*, employee PII could be *confidential*, and internal communication between employees through official channels might be *internal*.
- The process used to promote and demote data among categories.
- The access criteria for each data classification category.
- The kind of encryption key required for each category.

Environment classification

Your enterprise might have multiple environments, such as development, testing, sandbox, preproduction, and production. Each environment can contain different types of data and have different encryption requirements.

Your encryption framework should include the following information about your environments:

- Define your enterprise environments.
- Define the encryption requirements for each environment. For example, you might use a single encryption key for all the data categories in your development environment, and in your production environment, you might use different encryption keys for each business application or data classification category.

Change events and processes

Encryption standards are subject to frequent change so that you can keep up with the latest technologies, best practices, and innovations. The following are common change events that might initiate a revision of your encryption standards:

- Changes in the minimum length of encryption keys
- Changes in the strength of an encryption algorithm
- Changes to who can access encryption keys or how
- Changes to the rotation intervals for your keys
- Changes to the process for deleting keys
- Changes to the key storage location or policies
- Changes to the process for backing up and restoring keys

Your encryption framework should include the following to help prepare your organization to manage, implement, and communicate changes to the encryption standards or policy:

- **Change control process** – The purpose of this process is to plan and prepare for the upcoming change. When you need to change your encryption standards or policy, this repeatable and scalable process is designed to define:
 - How your organization assesses the impact of the change

- Who can initiate changes
- Who is responsible for implementing the change
- Who is responsible for approving the change
- How your organization would roll back the change, if necessary
- **Change auditability and traceability process** – This process defines how your organization audits and traces changes, both at the metadata level and at the data level. It should define how you keep and access records of:
 - What changed
 - When it was changed
 - Who initiated, approved, and implemented the change

For example, if your organization changes the minimum encryption key strength, you should be able to determine the original and new requirements, when the change was effective, and who was involved in the change process.

- **Change rollout process** – The purpose of this process is to define how your organization implements the change after you have decided to make it. This process defines:
 - Who are the stakeholders
 - Whether you should complete a pilot or proof of concept
 - How and when you should communicate the status of the change
 - How to roll back the change, if necessary.
 - What the observation period should be after implementing the change.
 - What the observation process will be to monitor the impact of the change, including how to collect feedback about the change and assess effectiveness
- **Retirement process** – The purpose of this process is to define how your organization handles retirement of encryption-related resources and information. It includes instructions for the actual retirement as well as the communication process for retirement.

Implementation

In this strategy, *architecture* refers to the technical implementation of your encryption standards. This section includes information about how AWS services, such as [AWS Key Management Service \(AWS KMS\)](#) and [AWS CloudHSM](#), can help you implement your data-at-rest encryption strategy according to your policy and standards.

AWS KMS is a managed service that helps you create and control the cryptographic keys that are used to protect your data. KMS keys never leave the service unencrypted. To use or manage your KMS keys, you interact with AWS KMS, and many AWS services are integrated with AWS KMS.

AWS CloudHSM is a cryptographic service for creating and maintaining *hardware security modules* (HSMs) in your AWS environment. HSMs are computing devices that process cryptographic operations and provide secure storage for cryptographic keys. If your standards requires you to use FIPS 140-2 Level 3 validated hardware, or if your standards dictate use of industry-standard APIs, such as PKCS#11, Java Cryptography Extensions (JCE), and Microsoft CryptoNG (CNG), then you might consider using AWS CloudHSM.

You can configure AWS CloudHSM as a custom key store for AWS KMS. This solution combines the convenience and service integration of AWS KMS with the added control and compliance benefits of using a AWS CloudHSM cluster in your AWS account. For more information, see [Custom key stores](#) (AWS KMS documentation).

This document discusses AWS KMS features at a high level and explains how AWS KMS can address your policy and standards.

Cost, convenience, and control

AWS KMS offers different types of keys. Some are owned or managed by AWS, and others are created and managed by customers. You can choose between these options based on the level of control you want to have over the key and cost considerations:

- **AWS owned keys** – AWS owns and manages these keys, and they are used in multiple AWS accounts. Some AWS services support AWS owned keys. You can use these keys for no charge. This key type relieves you from the cost and administrative overhead of managing the key lifecycle and access to it. For more information about this type of key, see [AWS owned keys](#) (AWS KMS documentation).
- **AWS managed keys** – If an AWS service is integrated with AWS KMS, it can create, manage, and use this type of keys on your behalf, in order to protect your resources in that service. These keys are created in your AWS account, and only AWS services can use them. There is no monthly fee for an AWS managed key. They can be subject to fees for use in excess of the free tier, but some AWS services cover these costs for you. You can use identity policies to control view and audit access for these keys, but AWS manages the key lifecycle. For more information about this type of key, see [AWS managed keys](#) (AWS KMS documentation). For a comprehensive list of the AWS services that integrate with AWS KMS, see [AWS service integration](#) (AWS marketing).

- **Customer managed keys** – You create, own, and manage this type of key, and you have full control over the key lifecycle. For segregation of duties, you can use both identity and resource-based policies to control access to the key. You can also set up automated [key rotation](#). Customer managed keys incur a monthly fee, and if you exceed the free tier, they also incur a fee for use. For more information about this type of key, see [Customer managed keys](#) (AWS KMS documentation).

For more information about key storage and usage, see [AWS Key Management Service pricing](#) (AWS marketing).

Performance and encryption types

Based upon the chosen encryption type in standards, you can use two types of KMS keys.

- **Symmetric** – All of the AWS KMS key types support symmetric encryption. When encrypting customer managed keys, you can use a single-strength key for encryption and decryption with AES-256-GCM.
- **Asymmetric** – Customer managed keys support asymmetric encryption. You can choose between different key strengths and algorithms, based upon your intended use. Asymmetric keys can encrypt and decrypt with RSA and can sign and verify operations with RSA or ECC. Asymmetric key algorithms inherently provide separation of roles and simplify key management. When using asymmetric encryption with AWS KMS, some operations aren't supported, such as rotating keys and importing external key material.

For more information about the AWS KMS operations that symmetric and asymmetric keys support, see [Key type reference](#) (AWS KMS documentation).

Envelope encryption

Envelope encryption is built into AWS KMS. In AWS KMS, you generate data keys in either plaintext or encrypted format. Encrypted data keys are encrypted with a KMS key. You can store the KMS key in a custom key store in an AWS CloudHSM cluster. For more information about the benefits of envelope encryption, see [About envelope encryption](#).

Key storage location

You use policies to manage access to AWS KMS resources. *Policies* describe who can access which resources. Policies attached to an AWS Identity and Access Management (IAM) principal are called

identity-based policies or *IAM policies*. Policies attached to other kinds of resources are called *resource policies*. AWS KMS resource policies for AWS KMS keys are called *key policies*. Every KMS key has a key policy.

Key policies provide flexibility to store the encryption key in a central location or store it closer to the data, in a distributed manner. Consider the following AWS KMS features when you're deciding where to store KMS keys in your AWS account:

- **Single-Region infrastructure support** – By default, KMS keys are Region-specific, and they never leave AWS KMS unencrypted. If your standards have strict requirements for controlling keys in a specific geographical location, explore using single-Region keys.
- **Multi-Region infrastructure support** – AWS KMS also supports special-purpose key type called *multi-Region keys*. Storing data in multiple AWS Regions is a common configuration for disaster recovery. By using multi-Region keys, you can transfer data between Regions without re-encrypting it, and you can manage the data as if you had the same key in each Region. This functionality is highly useful if your standards require that your encryption infrastructure spans multiple Regions in an active-active configuration. For more information, see [Multi-Region keys](#) (AWS KMS documentation).
- **Centralized management** – If your standards require that you store keys in a centralized location, you can use AWS KMS to store all of your encryption keys in a single AWS account. You use key policies to grant access to other applications, which can be in different accounts in the same Region. Centralized key management can reduce the administrative overhead of managing the key lifecycle and key access control.
- **External key material** – You can import externally generated key material into AWS KMS. Support for this functionality is available for single and multi-Region symmetric keys. Because the material of the symmetric key is generated externally, you are responsible for protecting the generated key materials. For more information, see [Imported key material](#) (AWS KMS documentation).

Access control

In AWS KMS, you can implement granular-level access control by using the following policy mechanisms: [key policies](#), [IAM policies](#), and [grants](#). Using these controls, you can set-up your separation of duties based on roles, such as administrators, key users who can encrypt the data, key users who can decrypt the data, and key users who can both encrypt and decrypt the data. For more information, see [Authentication and access control](#) (AWS KMS documentation).

Auditing and logging

AWS KMS integrates with AWS CloudTrail and Amazon EventBridge for logging and monitoring purposes. All AWS KMS API operations are recorded and auditable in CloudTrail logs. You can use Amazon CloudWatch, EventBridge, and AWS Lambda to set up custom monitoring solutions to configure notifications and automatic remediation. For more information, see [Logging and monitoring](#) (AWS KMS documentation).

FAQ

This section provides answers to commonly raised questions when defining your encryption standards or when creating your encryption infrastructure in the implementation phase.

When do I need symmetric encryption?

You might use symmetric encryption when:

- Speed, cost, and lower computational overhead are a priority.
- You need to encrypt a large amount of data.
- The encrypted data isn't leaving the boundaries of the organization's network.

When do I need asymmetric encryption?

You might use asymmetric encryption when:

- You need to share the data outside of the organization.
- Regulations or governance prohibit sharing the key.
- Nonrepudiation is required. (*Nonrepudiation* prevents a user from denying prior commitments or actions.)
- You need to strictly segregate access to encryption keys based on organization roles.

When do I need envelope encryption?

You need to support and implement envelope encryption if your encryption policy requires key rotation. Some governance and compliance regimes require key rotation, or your policy might mandate it to meet a business need.

When do I need to use a hardware security module (HSM)?

You might need an HSM if your policy specifies compliance with:

- The Federal Information Processing Standards (FIPS) 140-2 level 3 encryption standard. For more information, see [FIPS validation](#) (AWS CloudHSM documentation).

- Industry-standard APIs, such as PKCS#11, Java Cryptography Extension (JCE), or Microsoft Cryptography API: Next Generation (CNG)

Why should I centrally manage encryption keys?

The following are common benefits of centralized key management:

- Because keys are used and administered in different locations, you can reuse keys, which can reduce costs.
- You have more control over access to the encryption keys.
- Storing keys in a single location makes it easier to view, audit, and update keys in the event of a standards change.

Do I need to use a purpose-built encryption infrastructure for data at rest?

Your enterprise needs an encryption infrastructure if any one of the following is true:

- Your enterprise handles and stores data of any classification other than public.
- Your enterprise captures and stores data about employees or customers.
- Your enterprise handles PII data.
- Your enterprise must be compliant with regulatory or governance regimes that require data to be encrypted.
- Your enterprise executive leadership has mandated encryption of all data at rest.

How can AWS KMS help my organization meet its encryption objectives for data at rest?

In addition to many other features, AWS Key Management Service can help you:

- Use envelope encryption.
- Control encryption key access, such as separating key administration from key usage.
- Share keys across multiple AWS Regions and AWS accounts.

- Centralize key administration.
- Automate and mandate key rotation.

Resources

AWS service documentation

- [AWS KMS Cryptographic Details](#)
- [AWS KMS Developer Guide](#)
 - [AWS KMS concepts](#)
 - [Special-purpose keys](#)
 - [Authentication and access control for AWS KMS](#)
 - [Security of AWS KMS](#)
 - [How AWS services use AWS KMS](#)
- [AWS CloudHSM User Guide](#)
- [AWS cryptographic services and tools guide](#)
 - [How to choose an encryption tool or service](#)
 - [Cryptography concepts](#)

AWS marketing

- [AWS KMS pricing](#)
- [AWS KMS integration with other AWS services](#)

AWS Well-Architected Framework

- [Protecting data in transit](#)
- [Protecting data at rest](#)

Hashing and tokenization

- [How to use tokenization to improve data security and reduce audit scope](#) (AWS blog post)
- [Recommendation for applications using approved hash algorithms](#) (NIST publication)

Videos

- [How encryption works in AWS](#)
- [Securing your block storage on AWS](#)
- [Achieving security goals with AWS CloudHSM](#)
- [Best practices for implementing AWS Key Management Service](#)
- [A deep dive into AWS encryption services](#)

Document history

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

Change	Description	Date
Initial publication	—	September 15, 2022

AWS Prescriptive Guidance glossary

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

Numbers

7 Rs

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

- Refactor/re-architect – Move an application and modify its architecture by taking full advantage of cloud-native features to improve agility, performance, and scalability. This typically involves porting the operating system and database. Example: Migrate your on-premises Oracle database to the Amazon Aurora PostgreSQL-Compatible Edition.
- Replatform (lift and reshape) – Move an application to the cloud, and introduce some level of optimization to take advantage of cloud capabilities. Example: Migrate your on-premises Oracle database to Amazon Relational Database Service (Amazon RDS) for Oracle in the AWS Cloud.
- Repurchase (drop and shop) – Switch to a different product, typically by moving from a traditional license to a SaaS model. Example: Migrate your customer relationship management (CRM) system to Salesforce.com.
- Rehost (lift and shift) – Move an application to the cloud without making any changes to take advantage of cloud capabilities. Example: Migrate your on-premises Oracle database to Oracle on an EC2 instance in the AWS Cloud.
- Relocate (hypervisor-level lift and shift) – Move infrastructure to the cloud without purchasing new hardware, rewriting applications, or modifying your existing operations. 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.

A

ABAC

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

abstracted services

See [managed services](#).

ACID

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

active-active migration

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

active-passive migration

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

aggregate function

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

AI

See [artificial intelligence](#).

AIOps

See [artificial intelligence operations](#).

anonymization

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

anti-pattern

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

application control

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

application portfolio

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

artificial intelligence (AI)

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

artificial intelligence operations (AIOps)

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

asymmetric encryption

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

atomicity, consistency, isolation, durability (ACID)

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

attribute-based access control (ABAC)

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

authoritative data source

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

Availability Zone

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

AWS Cloud Adoption Framework (AWS CAF)

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

AWS Workload Qualification Framework (AWS WQF)

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

B

bad bot

A [bot](#) that is intended to disrupt or cause harm to individuals or organizations.

BCP

See [business continuity planning](#).

behavior graph

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

big-endian system

A system that stores the most significant byte first. See also [endianness](#).

binary classification

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

bloom filter

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

blue/green deployment

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

bot

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

botnet

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

branch

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

break-glass access

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

brownfield strategy

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

buffer cache

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

business capability

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

business continuity planning (BCP)

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

C

CAF

See [AWS Cloud Adoption Framework](#).

canary deployment

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

CCoE

See [Cloud Center of Excellence](#).

CDC

See [change data capture](#).

change data capture (CDC)

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

chaos engineering

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

CI/CD

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

classification

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

client-side encryption

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

Cloud Center of Excellence (CCoE)

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

cloud computing

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

cloud operating model

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

cloud stages of adoption

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

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

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

CMDB

See [configuration management database](#).

code repository

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

cold cache

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

cold data

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

computer vision (CV)

A field of [AI](#) that uses machine learning to analyze and extract information from visual formats such as digital images and videos. For example, AWS Panorama offers devices that add CV to on-premises camera networks, and Amazon SageMaker provides image processing algorithms for CV.

configuration drift

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

configuration management database (CMDB)

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

conformance pack

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

continuous integration and continuous delivery (CI/CD)

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

CV

See [computer vision](#).

D

data at rest

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

data classification

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

data drift

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

data in transit

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

data mesh

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

data minimization

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

data perimeter

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

data preprocessing

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

data provenance

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

data subject

An individual whose data is being collected and processed.

data warehouse

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

database definition language (DDL)

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

database manipulation language (DML)

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

DDL

See [database definition language](#).

deep ensemble

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

deep learning

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

defense-in-depth

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

delegated administrator

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

deployment

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

development environment

See [environment](#).

detective control

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

development value stream mapping (DVSM)

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

digital twin

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

dimension table

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

disaster

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

disaster recovery (DR)

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

DML

See [database manipulation language](#).

domain-driven design

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

DR

See [disaster recovery](#).

drift detection

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

DVSM

See [development value stream mapping](#).

E

EDA

See [exploratory data analysis](#).

edge computing

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

encryption

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

encryption key

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

endianness

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

endpoint

See [service endpoint](#).

endpoint service

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

enterprise resource planning (ERP)

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

envelope encryption

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

environment

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

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

epic

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

ERP

See [enterprise resource planning](#).

exploratory data analysis (EDA)

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

F

fact table

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

fail fast

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

fault isolation boundary

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

feature branch

See [branch](#).

features

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

feature importance

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

feature transformation

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

FGAC

See [fine-grained access control](#).

fine-grained access control (FGAC)

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

flash-cut migration

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

G

geo blocking

See [geographic restrictions](#).

geographic restrictions (geo blocking)

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

Gitflow workflow

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

greenfield strategy

The absence of existing infrastructure in a new environment. When adopting a greenfield strategy for a system architecture, you can select all new technologies without the restriction of compatibility with existing infrastructure, also known as [brownfield](#). If you are expanding the existing infrastructure, you might blend brownfield and greenfield strategies.

guardrail

A high-level rule that helps govern resources, policies, and compliance across organizational units (OUs). *Preventive guardrails* enforce policies to ensure alignment to compliance standards. They are implemented by using service control policies and IAM permissions boundaries. *Detective guardrails* detect policy violations and compliance issues, and generate alerts

for remediation. They are implemented by using AWS Config, AWS Security Hub, Amazon GuardDuty, AWS Trusted Advisor, Amazon Inspector, and custom AWS Lambda checks.

H

HA

See [high availability](#).

heterogeneous database migration

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

high availability (HA)

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

historian modernization

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

homogeneous database migration

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

hot data

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

hotfix

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

hypercare period

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

I

laC

See [infrastructure as code](#).

identity-based policy

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

idle application

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

IIoT

See [industrial Internet of Things](#).

immutable infrastructure

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

inbound (ingress) VPC

In an AWS multi-account architecture, a VPC that accepts, inspects, and routes network connections from outside an application. The [AWS Security Reference Architecture](#) recommends

setting up your Network account with inbound, outbound, and inspection VPCs to protect the two-way interface between your application and the broader internet.

incremental migration

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

Industry 4.0

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

infrastructure

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

infrastructure as code (IaC)

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

industrial Internet of Things (IIoT)

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

inspection VPC

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

Internet of Things (IoT)

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

interpretability

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

IoT

See [Internet of Things.](#)

IT information library (ITIL)

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

IT service management (ITSM)

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

ITIL

See [IT information library.](#)

ITSM

See [IT service management.](#)

L

label-based access control (LBAC)

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

landing zone

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

large migration

A migration of 300 or more servers.

LBAC

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

least privilege

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

lift and shift

See [7 Rs](#).

little-endian system

A system that stores the least significant byte first. See also [endianness](#).

lower environments

See [environment](#).

M

machine learning (ML)

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

main branch

See [branch](#).

malware

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

managed services

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

manufacturing execution system (MES)

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

MAP

See [Migration Acceleration Program](#).

mechanism

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

member account

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

MES

See [manufacturing execution system](#).

Message Queuing Telemetry Transport (MQTT)

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

microservice

A small, independent service that communicates over well-defined APIs and is typically owned by small, self-contained teams. For example, an insurance system might include

microservices that map to business capabilities, such as sales or marketing, or subdomains, such as purchasing, claims, or analytics. The benefits of microservices include agility, flexible scaling, easy deployment, reusable code, and resilience. For more information, see [Integrating microservices by using AWS serverless services](#).

microservices architecture

An approach to building an application with independent components that run each application process as a microservice. These microservices communicate through a well-defined interface by using lightweight APIs. Each microservice in this architecture can be updated, deployed, and scaled to meet demand for specific functions of an application. For more information, see [Implementing microservices on AWS](#).

Migration Acceleration Program (MAP)

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

migration at scale

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

migration factory

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

migration metadata

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

migration pattern

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

Migration Portfolio Assessment (MPA)

An online tool that provides information for validating the business case for migrating to the AWS Cloud. MPA provides detailed portfolio assessment (server right-sizing, pricing, TCO comparisons, migration cost analysis) as well as migration planning (application data analysis and data collection, application grouping, migration prioritization, and wave planning). The [MPA tool](#) (requires login) is available free of charge to all AWS consultants and APN Partner consultants.

Migration Readiness Assessment (MRA)

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

migration strategy

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

ML

See [machine learning](#).

modernization

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

modernization readiness assessment

An evaluation that helps determine the modernization readiness of an organization's applications; identifies benefits, risks, and dependencies; and determines how well the organization can support the future state of those applications. The outcome of the assessment is a blueprint of the target architecture, a roadmap that details development phases and

milestones for the modernization process, and an action plan for addressing identified gaps. For more information, see [Evaluating modernization readiness for applications in the AWS Cloud](#).

monolithic applications (monoliths)

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

MPA

See [Migration Portfolio Assessment](#).

MQTT

See [Message Queuing Telemetry Transport](#).

multiclass classification

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

mutable infrastructure

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

O

OAC

See [origin access control](#).

OAI

See [origin access identity](#).

OCM

See [organizational change management](#).

offline migration

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

OI

See [operations integration](#).

OLA

See [operational-level agreement](#).

online migration

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

OPC-UA

See [Open Process Communications - Unified Architecture](#).

Open Process Communications - Unified Architecture (OPC-UA)

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

operational-level agreement (OLA)

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

operational readiness review (ORR)

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

operational technology (OT)

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

operations integration (OI)

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

organization trail

A trail that's created by AWS CloudTrail that logs all events for all AWS accounts in an organization in AWS Organizations. This trail is created in each AWS account that's part of the organization and tracks the activity in each account. For more information, see [Creating a trail for an organization](#) in the CloudTrail documentation.

organizational change management (OCM)

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

origin access control (OAC)

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

origin access identity (OAI)

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

ORR

See [operational readiness review](#).

OT

See [operational technology](#).

outbound (egress) VPC

In an AWS multi-account architecture, a VPC that handles network connections that are initiated from within an application. The [AWS Security Reference Architecture](#) recommends

setting up your Network account with inbound, outbound, and inspection VPCs to protect the two-way interface between your application and the broader internet.

P

permissions boundary

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

personally identifiable information (PII)

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

PII

See [personally identifiable information](#).

playbook

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

PLC

See [programmable logic controller](#).

PLM

See [product lifecycle management](#).

policy

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

polyglot persistence

Independently choosing a microservice's data storage technology based on data access patterns and other requirements. If your microservices have the same data storage technology, they can encounter implementation challenges or experience poor performance. Microservices are more easily implemented and achieve better performance and scalability if they use the data store best adapted to their requirements. For more information, see [Enabling data persistence in microservices](#).

portfolio assessment

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

predicate

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

predicate pushdown

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

preventative control

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

principal

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

Privacy by Design

An approach in system engineering that takes privacy into account throughout the whole engineering process.

private hosted zones

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

proactive control

A [security control](#) designed to prevent the deployment of noncompliant resources. These controls scan resources before they are provisioned. If the resource is not compliant with the control, then it isn't provisioned. For more information, see the [Controls reference guide](#) in the AWS Control Tower documentation and see [Proactive controls](#) in *Implementing security controls on AWS*.

product lifecycle management (PLM)

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

production environment

See [environment](#).

programmable logic controller (PLC)

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

pseudonymization

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

publish/subscribe (pub/sub)

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

Q

query plan

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

query plan regression

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

R

RACI matrix

See [responsible, accountable, consulted, informed \(RACI\)](#).

ransomware

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

RASCI matrix

See [responsible, accountable, consulted, informed \(RACI\)](#).

RCAC

See [row and column access control](#).

read replica

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

re-architect

See [7 Rs](#).

recovery point objective (RPO)

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

recovery time objective (RTO)

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

refactor

See [7 Rs](#).

Region

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

regression

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

rehost

See [7 Rs](#).

release

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

relocate

See [7 Rs](#).

replatform

See [7 Rs](#).

repurchase

See [7 Rs](#).

resiliency

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

resource-based policy

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

responsible, accountable, consulted, informed (RACI) matrix

A matrix that defines the roles and responsibilities for all parties involved in migration activities and cloud operations. The matrix name is derived from the responsibility types defined in the

matrix: responsible (R), accountable (A), consulted (C), and informed (I). The support (S) type is optional. If you include support, the matrix is called a *RASCI matrix*, and if you exclude it, it's called a *RACI matrix*.

responsive control

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

retain

See [7 Rs](#).

retire

See [7 Rs](#).

rotation

The process of periodically updating a [secret](#) to make it more difficult for an attacker to access the credentials.

row and column access control (RCAC)

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

RPO

See [recovery point objective](#).

RTO

See [recovery time objective](#).

runbook

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

S

SAML 2.0

An open standard that many identity providers (IdPs) use. This feature enables federated single sign-on (SSO), so users can log into the AWS Management Console or call the AWS API

operations without you having to create user in IAM for everyone in your organization. For more information about SAML 2.0-based federation, see [About SAML 2.0-based federation](#) in the IAM documentation.

SCADA

See [supervisory control and data acquisition](#).

SCP

See [service control policy](#).

secret

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

security control

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

security hardening

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

security information and event management (SIEM) system

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

security response automation

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

server-side encryption

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

service control policy (SCP)

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

service endpoint

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

service-level agreement (SLA)

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

service-level indicator (SLI)

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

service-level objective (SLO)

A target metric that represents the health of a service, as measured by a [service-level indicator](#).

shared responsibility model

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

SIEM

See [security information and event management system](#).

single point of failure (SPOF)

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

SLA

See [service-level agreement](#).

SLI

See [service-level indicator](#).

SLO

See [service-level objective](#).

split-and-seed model

A pattern for scaling and accelerating modernization projects. As new features and product releases are defined, the core team splits up to create new product teams. This helps scale your organization's capabilities and services, improves developer productivity, and supports rapid innovation. For more information, see [Phased approach to modernizing applications in the AWS Cloud](#).

SPOF

See [single point of failure](#).

star schema

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

strangler fig pattern

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

subnet

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

supervisory control and data acquisition (SCADA)

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

symmetric encryption

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

synthetic testing

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

T

tags

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

target variable

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

task list

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

test environment

See [environment](#).

training

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

transit gateway

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

trunk-based workflow

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

trusted access

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

tuning

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

two-pizza team

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

U

uncertainty

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

undifferentiated tasks

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

upper environments

See [environment](#).

V

vacuuming

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

version control

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

VPC peering

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

vulnerability

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

W

warm cache

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

warm data

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

window function

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

workload

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

workstream

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

WORM

See [write once, read many](#).

WQF

See [AWS Workload Qualification Framework](#).

write once, read many (WORM)

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

Z

zero-day exploit

An attack, typically malware, that takes advantage of a [zero-day vulnerability](#).

zero-day vulnerability

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

zombie application

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