Designing an AWS Control Tower landing zone

AWS Prescriptive Guidance

AWS Prescriptive Guidance: Designing an AWS Control Tower landing zone

Table of Contents

Designing an AWS Control Tower landing zone

Vikas Dewangan, Emelie Akerstrom, and Pooja Banerjee, Amazon Web Services (AWS)

December 2024 (document history)

Typically, you begin a cloud adoption journey by conceptualizing and designing a landing zone, which is a well-architected, multi-account AWS environment that is scalable and secure. A landing zone creates an agile and scalable cloud environment. It also helps you quickly launch and deploy workloads and applications into your infrastructure, which helps accelerate your digital transformation and cloud journey. You can use AWS Control Tower to set up and govern a secure, multi-account AWS environment for your landing zone based on AWS best practices. The landing zone includes foundational AWS services that are created before you deploy workloads.

Enterprises can use a landing zone at various stages of their cloud adoption journey. Some enterprises have recently started their journey on AWS, aren't aware of AWS best practices for multi-account strategy, security, and networking, and aren't sure how to set up their landing zones. Other enterprises have been using AWS for a while and want to scale their landing zone to deploy additional workloads.

This guide is for enterprises that have decided to set up an AWS Control Tower landing zone and want to create a design document that covers the key functionality of the landing zone. The guide covers best practices and patterns for setting up the account structure and configuring networking, logging, authentication, and other aspects of the landing zone. It also provides examples of how to present this information in your design documents. The key outcome from this exercise is a design artifact that documents the decisions that were made and the architectures that were agreed upon as the basis of implementing the landing zone. This could be followed by new cloud-native application development or the migration of existing workloads into the landing zone.

Design process

IT infrastructure teams and enterprise architects usually create a design document that addresses the important parts of a landing zone. Stakeholders must approve the design document before the landing zone is implemented. This guide accelerates the design process for IT infrastructure teams, enterprise architects, and cloud migration teams by providing a landing zone design template that's aligned with the AWS Well-Architected Framework. It helps you accelerate your cloud adoption by outlining the foundational pillars of the landing zone design. The key landing zone topics covered in this design guide are:

- Setting up a landing zone – Provides an overview of the basic landing zone setup, including where it should be deployed, which parameters to use, and which resources are deployed.
- Configuring account structure and OUs – Shares a sample account structure that includes different accounts and organizational units (OUs) for a landing zone.
- Using controls to govern resources and monitor compliance – Identifies and explains the preventive, detective, and proactive controls that must be enabled in AWS Control Tower for the landing zone's governance.
- Networking integration – Helps you design connectivity between the virtual private clouds (VPCs) in your landing zone and on-premises applications by using native AWS services such as AWS Transit Gateway, AWS Direct Connect, and AWS Site-to-Site VPN.
- Authentication and authorization – Explains how AWS IAM Identity Center integrates with AWS Control Tower and how to integrate your own identity provider.
- Centralizing logging and monitoring – Reviews the monitoring, logging, log archival, and alerting strategy for the landing zone.
- Managing the configuration of AWS resources – Explains how you can manage resource configurations, track changes, and view compliance data by using AWS Config.

AWS services

This design guide covers the following AWS services to set up the landing zone:

- AWS Control Tower automates the landing zone setup by using best practices for identity, federated access, controls, and account structure.
- AWS CloudTrail integrates with AWS Control Tower, captures actions as events, and provides a record of actions taken by a user, role, or AWS service.
- AWS Config helps you assess, audit, and evaluate the configurations of your AWS resources. AWS Control Tower automatically enables AWS Config in its AWS Regions and uses it to implement detective controls. The configuration history and snapshots are delivered to an Amazon Simple Storage Service (Amazon S3) bucket in the Log Archive account. (See Security OU – Log Archive account in the *AWS Security Reference Architecture*.)
- AWS Direct Connect establishes a dedicated network connection from on-premises environments to the AWS Cloud.
- AWS Identity and Access Management (IAM) helps securely control access to your AWS resources. IAM or IAM Identity Center authenticates that you're an approved user before you perform

operations, such as provisioning accounts in AWS Control Tower Account Factory or creating new OUs in the AWS Control Tower console, in your landing zone.

- AWS Organizations is an account management service that consolidates multiple AWS accounts into an organization that you centrally manage. In AWS Control Tower, AWS Organizations helps you to centrally manage billing, control access, compliance, and security, in addition to sharing resources across your member accounts.
- Amazon CloudWatch monitors the resources and applications that run on AWS. CloudWatch cross-account observability lets you monitor and troubleshoot applications that span multiple accounts within a Region.
- Amazon GuardDuty continuously monitors for malicious activity and unauthorized behavior to protect your AWS accounts, workloads, and the data stored in Amazon S3.
- AWS IAM Identity Center is a cloud-based single sign-on (SSO) service that helps you centrally manage SSO access for all your AWS accounts and cloud applications. AWS Control Tower integrates with IAM Identity Center to manage users, roles, and multi-account access.
- Amazon Simple Storage Service (Amazon S3) is a cloud-based object storage service that helps you store, protect, and retrieve any amount of data. S3 buckets store the AWS Control Tower logs and AWS access logs.
- AWS Systems Manager provides a unified user interface (UI) to centrally track and resolve operational issues across your applications and resources.
- Amazon Simple Notification Service (Amazon SNS) helps you coordinate and manage the exchange of messages between publishers and clients, including web servers and email addresses. AWS Control Tower sends security notifications, such as non-compliance with detective controls, to SNS topics that clients can subscribe to.

Targeted business outcomes

You should expect the following business outcomes from using this design guide:

- Lead design conversations for an AWS Control Tower landing zone confidently.
- Address all foundational pillars of an AWS Control Tower landing zone effectively.
- Create a strong foundation for migrating workloads to the AWS Cloud.

Setting up a landing zone

AWS Control Tower automates the setup of a landing zone by using best-practice templates for identity, federated access, controls, and account structure. AWS Control Tower offers the easiest way to set up and govern a secure, multi-account AWS environment. From a feature perspective, the AWS Control Tower dashboard provides visibility into your landing zone environment, an aggregated view of organizational units (OUs) and accounts in the organization, the controls that are enabled, and the compliance status of OUs and accounts with those controls. A list of non-compliant resources is also provided to identify any required actions.

If your enterprise is new to AWS, we recommend that you start with AWS Control Tower as the foundation for your landing zone. However, as an alternative, you can opt for a custom-built landing zone.

AWS Control Tower managed resources

When you set up a landing zone, AWS Control Tower creates multiple managed resources in your management account for the accounts in the landing zone. For more information about these resources, see How AWS Control Tower works in the AWS Control Tower documentation.

> **Important**
>
> When you use a landing zone, AWS Control Tower managed resources must not be modified or deleted; otherwise, the landing zone might drift and enter an unknown state. For more information about drift, see Detect and resolve drift in AWS Control Tower in the AWS Control Tower documentation.

Deployment

AWS Control Tower is deployed in a management account that's created for the new landing zone. The management account is where you provision new accounts and centrally manage or configure controls, user access, permissions, and OUs.

When you set up the landing zone, AWS Control Tower automatically creates a Security OU that contains the Log Archive account and Audit account. These accounts enable centralized management and governance of the landing zone through monitoring and logging. For more

information about OUs and accounts in the landing zone, see the Account structure and OUs section of this guide.

Launch parameters

The following table shows the required parameters for setting up a landing zone.

Category	Parameter	Sample values
AWS Regions	*Home Region* This is the AWS Region where shared resources will be provisioned. You cannot change the home Region after the landing zone is set up, but you can add more Regions to govern.	Your current Region (as reflected in the AWS Region selector on the navigation bar). This field is not editable.
	Additional Regions for governance You can optionally choose additional Regions for AWS Control Tower to govern.	No additional Regions
	Region deny setting If you have data residency requirements, you can optionally choose to enable a Region deny service control policy (SCP) to deny access in Regions that aren't selected.	Not enabled
Configuration for the landing zone shared accounts	*Foundational OU* This is the OU that contains the shared accounts.	Security

Category	Parameter	Sample values
	Additional OU This is a secondary OU for storing production or development accounts.	Sandbox
	Log Archive account options Either create a new account or use your existing Log Archive account.	Create new account
	Log Archive account details If you are creating a new account, specify a unique email address that is not yet associated with an AWS account. You can also specify the account name. The default name is Log Archive. If you are using an existing Log Archive account, specify the AWS account ID. **Note**: These details cannot be edited after the landing zone has been set up.	aws-logarchive@example.com Log Archive
	Audit account options Either create a new account or use your existing Audit account.	Create new account

Category	Parameter	Sample values
	Audit account details	audit@example.com
	If you are creating a new account, specify a unique email address that is not yet associated with an AWS account. Also specify the account name. The default name is Audit. If you are using an existing audit account, specify the AWS account ID. **Note:** These details cannot be edited after the landing zone has been set up.	Audit
Additional configurations **Note:** These are optional configurations. You can leave them at their default settings or choose your own configura tion.	*AWS account access configura tion* You can optionally choose to manage account access yourself or accept the default IAM Identity Center setup in AWS Control Tower.	AWS Control Tower sets up AWS account access with IAM Identity Center.
	AWS CloudTrail configuration You can optionally choose to manage CloudTrail in your organization yourself or accept the default CloudTrai l setup from AWS Control Tower. The default setting enables an organization-level trail for management events in your Log Archive account.	Enabled

Category	Parameter	Sample values
	Log configuration for Amazon S3 You can optionally configure log retention for the Log Archive S3 bucket or accept the default retention periods.	Standard account logging: 1 year Access logging: 10 years
	KMS encryption You can optionally enable encryption for AWS Control Tower resources by using an AWS Key Management Service (AWS KMS) customer managed key. If you enable encryption, you are asked to specify the key name or Amazon Resource Name (ARN) of the customer managed key to be used. **Note:** If you don't enable this option, AWS Control Tower uses SSE-S3 encryption with AWS managed keys as the default configuration.	Disabled

Configuring account structure and OUs

You can achieve both speed and security in the cloud by using AWS accounts. AWS accounts are resource containers that assist users in managing AWS resources. As needs and demands grow, environments scale through the addition of new accounts. The use of multiple accounts is a best practice because it enables rapid innovation across distributed teams and individuals, reduces the scope of impact through isolation, and can be adjusted to meet new business processes based on operational, regulatory, and budgetary requirements. Lastly, costs are incurred at the account level, so activity and costs can be identified with each account.

One of the key principles of account structure design is to start with a basic structure and expand it as your needs evolve. For an example of an incremental approach for account structure design, see the section Patterns for organizing your AWS accounts in the AWS whitepaper *Organizing Your AWS Environment Using Multiple Accounts*. Your account structure might have to be tailored based on your industry. For examples, see the AWS blog post Defining an AWS multi-account strategy for a digital bank and the AWS Prescriptive Guidance guide OU structure in regulated AWS landing zones: an example from the pharmaceutical industry. Let's look at some key considerations for OU design.

- **Preventive and detective controls**. A key consideration for defining a separate OU for a group of accounts is whether there are a set of controls you would like to apply to that OU. If you are planning to use custom controls, specifically SCPs, consider the quota on the number of SCPs that you can attach to an OU. If you need more custom SCPs for an OU than the quota permits, you can design your account structure to apply some of those SCPs at an OU that's one level higher or lower to achieve the desired outcome.
- **Shared resources**. Another consideration is whether there are any AWS resources (such as transit gateways or Amazon Route 53 forwarding rules) that you would like to share with all the accounts in the OU by using AWS Resource Access Manager.
- **Automation**. This can be an important consideration for OU design. For example, AWS CloudFormation StackSets supports deploying resources to all accounts within an OU, so you can group accounts that have common AWS resource requirements in an OU to automate deployments.

The following diagram provides an example of an AWS Control Tower–based account structure that includes various accounts and OUs and follows AWS best practices. You can customize this architecture based on the account structure that's suitable for your enterprise.

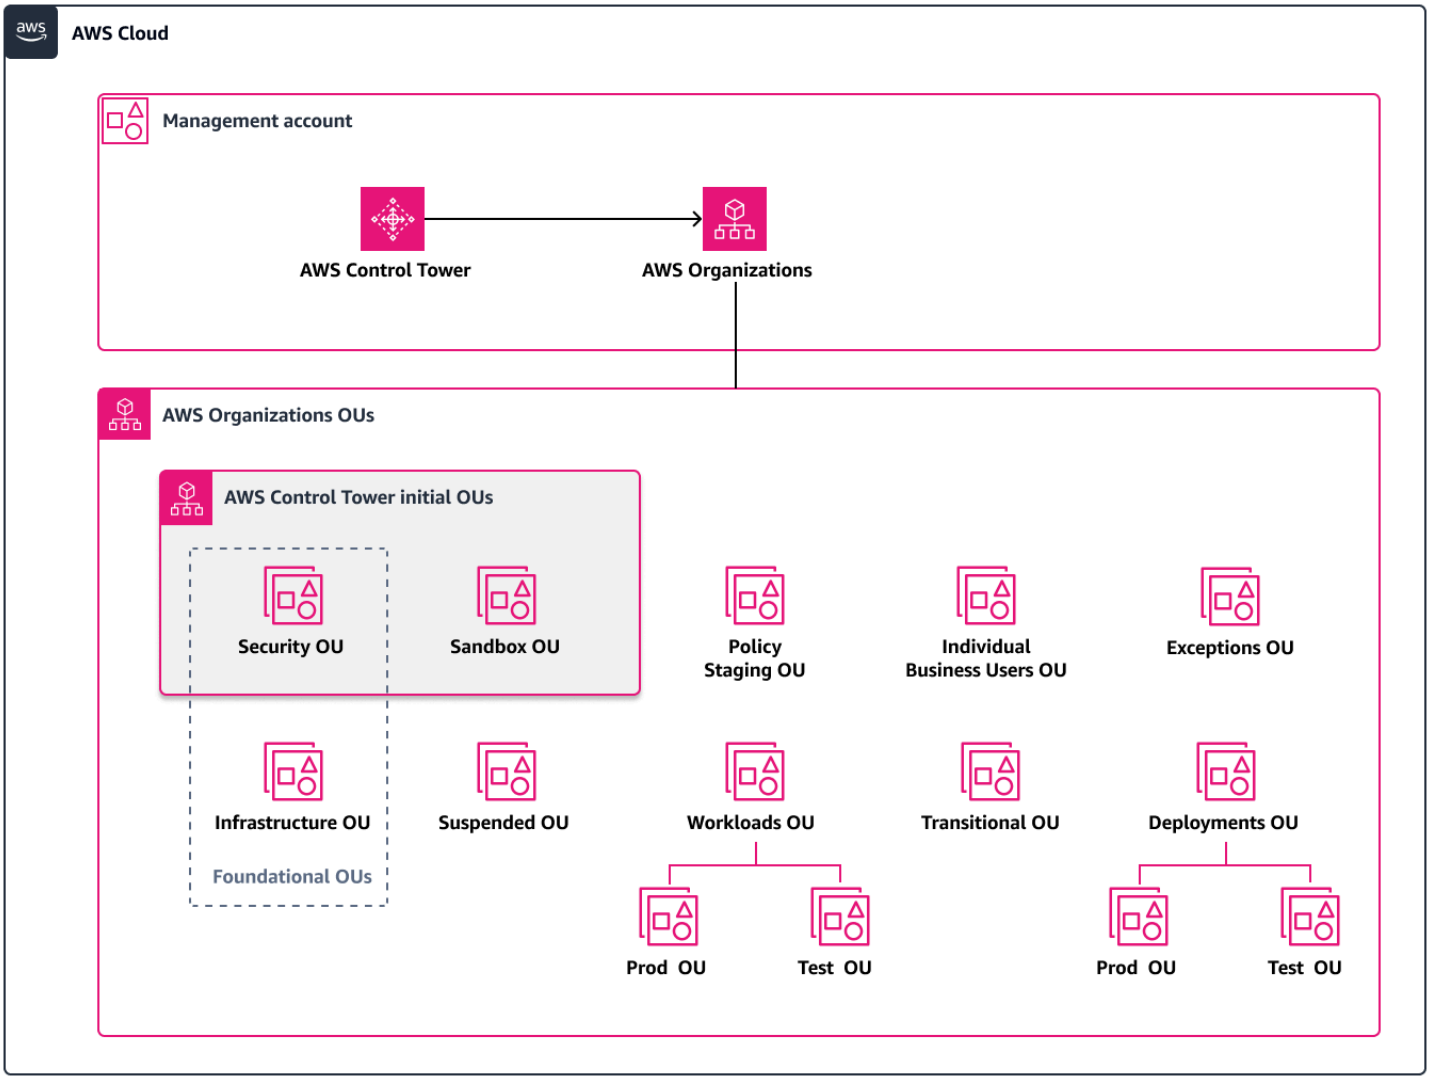

The account in which AWS Control Tower is deployed is automatically configured as the AWS Control Tower management account. The management account consolidates billing of all accounts in the landing zone. This account is used to provision new AWS accounts with AWS Control Tower Account Factory, manage OUs and controls, and manage user access and permissions by using IAM Identity Center.

AWS Control Tower sets up the Security OU that contains the Log Archive and Audit accounts, and provides an option to create an additional OU. You can change the names of these accounts during AWS Control Tower setup but not later. If required, you can specify an existing AWS account as the Log Archive or Audit account during the setup process. You have to provision all additional OUs and member accounts. You might decide to start with a subset of OUs at their beginning of your landing zone implementation and add OUs later.

The diagram includes the following key OUs.

OU	Description
Security OU	Automatically set up by AWS Control Tower to host the Log Archive and Audit accounts. • **Log Archive account** – AWS Control Tower centralizes all AWS CloudTrail and AWS Config logs into a centralized S3 bucket. You can also centralize other logs from across your organization, such as Amazon CloudWatch logs, Amazon S3 access logs, and VPC Flow Logs, in this account. • **Audit account** – Provides read and write access to all accounts in your landing zone. From the Audit account, you can programmatically review accounts by using a role that is granted only to AWS Lambda functions. Your security and compliance teams can use the Audit account to audit and review the accounts in your organization.
Sandbox OU	Set up by AWS Control Tower to host sandbox accounts that help you safely test and develop new services, processes, and templates.
Infrastructure OU	Contains the following accounts that are required for shared infrastructure services for your production and non-production workloads: • Backup account • Identity account • Monitoring account • Network account • Operations Tooling account • Shared Services account

OU	Description
	For details of the recommended AWS services and the AWS solutions that you can host in each account, see Infrastructure OU in the AWS whitepaper *Organizing Your AWS Environment Using Multiple Accounts*.
Policy Staging OU	Helps you safely test policy changes such as detective and preventive controls before applying them to OUs or accounts.
Individual Business Users OU	Hosts accounts for individual business users and teams who want to internally use AWS resources that aren't classified as workloads.
Exceptions OU	Hosts accounts that require exceptions to the default security policies that are applied to the Workloads OU.
Workloads OU	Hosts business workloads for both production and non-production environments.
Deployments OU	Hosts resources for building, validating, and promoting releases to your environments. Includes continuous integration and continuou s delivery (CI/CD) tooling.
Transitional OU	Provides a holding area for accounts that are being moved to the landing zone before formally integrating them into the standardi zed OUs in the account structure.
Suspended OU	Provides a locked and extremely restricte d environment to host suspended, deleted, and reused accounts. This is useful if you suspect that an account has been breached or compromised.

For further guidance on defining a multi-account strategy, see the AWS whitepaper Organizing Your AWS Environment Using Multiple Accounts.

You can automatically provision new AWS accounts by using AWS Control Tower Account Factory or AWS Control Tower Account Factory for Terraform (AFT), as described in the following sections.

AWS Control Tower Account Factory

AWS Control Tower Account Factory is a console-based AWS Control Tower feature that you can use to provision new accounts in your organization. Account Factory functions as a UI for an AWS Service Catalog product and provisions custom accounts by using AWS CloudFormation. You can also configure Account Factory to optionally create VPCs and a NAT gateway in the new accounts. This configuration automatically provisions a VPC with up to two subnets and a defined Classless InterDomain Routing (CIDR) range in the AWS Regions that you specify.

> **Note**
>
> The subnet IP address ranges in the VPCs of the accounts provisioned by Account Factory might overlap.

AWS Control Tower add-ons

AWS Control Tower add-on solutions deliver enhanced features and customizations provided by AWS to augment the capabilities of AWS Control Tower. These enhancements are designed to seamlessly elevate the functionality of AWS Control Tower to help meet organizational requirements for increased efficiency, security, and compliance within the AWS environment. Examples include the following:

- Landing Zone Accelerator (LZA) is an open source solution from AWS that's designed to expedite the setup of a secure, multi-account AWS environment based on AWS best practices. It streamlines the deployment of a landing zone and offers automated processes for account creation, configurable security and compliance guardrails, and customizations to align with your organization's policies.
- Account Factory for Terraform (AFT) is a Terraform module from AWS that extends the capabilities of the AWS Control Tower Account Factory. It enables organizations to create and manage AWS accounts by using infrastructure as code (IaC). AFT facilitates the definition of

account configurations and resources through version-controlled code, which helps ensure consistency and repeatability in the account creation process.

- Customizations for Control Tower (CfCT) is a solution by AWS that supports the deployment of resources by using CloudFormation templates and service control policies (SCPs). CfCT is integrated with AWS Control Tower lifecycle events and ensures that your deployments are synchronized with landing zone events. For example, it ensures that newly provisioned accounts are equipped with the correct infrastructure and automatically deploys domain-specific resources for accounts that are placed within OUs.

Using controls to govern resources and monitor compliance

AWS Control Tower controls are high-level rules that provide ongoing governance and enforce specific policies for your AWS environment. Controls can be applied to organizational units (OUs) and have three different types: *preventive*, *detective,* and *proactive*.

- **Preventive controls** help ensure that your accounts maintain compliance by disallowing actions that cause policy violations. Preventive controls are implemented with service control policies (SCPs), which are part of AWS Organizations. For example, the control Disallow Actions as a Root User helps ensure that the high privilege root user can't be used for unrestricted access to all resources in an account. Instead, users are forced to use more restricted IAM roles.
- **Detective controls** continuously monitor resources to detect non-compliance in your accounts, and then provide alerts through the dashboard. For example, the control Detect Whether Unrestricted Incoming TCP Traffic is Allowed can detect whether a security group is set up with unrestricted incoming TCP traffic and alert the user to restrict their incoming protocols. Detective controls are implemented by using AWS Config Rules and AWS Lambda functions.
- **Proactive controls** use AWS CloudFormation Hooks to help ensure that custom configuration and compliance checks are automatically enforced during the deployment of CloudFormation resources. These controls make it easier to maintain a secure and compliant AWS environment.

> **Note**
>
> SCPs (preventive controls) don't have any effect in the management account. The root user and IAM administrators in the management account can perform any action that is denied in an SCP. This ensures that the management account retains full administrative control over the organization and can't be accidentally locked out by any SCP errors. All actions that are performed in the management account are still tracked by the AWS CloudTrail and AWS Config recorder and stored in the Log Archive account.

Control guidance levels

AWS Control Tower controls have three different guidance levels: mandatory, strongly recommended, and elective.

Mandatory controls are automatically enabled and enforced by AWS Control Tower. Strongly recommended controls are optional and based on AWS best practices. Elective controls are also optional but are commonly used by enterprises. For more information, see the controls library in the AWS Control Tower documentation.

> **Note**
>
> You can use custom SCPs and AWS Config Rules for additional detection and prevention. These aren't implemented in AWS Control Tower but can be implemented in AWS Organizations and AWS Config.

Limitations for preventive controls

You can have a maximum of five SCPs attached to an OU and a maximum of five OU levels. This includes both custom SCPs and AWS Control Tower–created SCPs, so try to consolidate your SCPs into fewer documents. (AWS Control Tower will do this automatically for its preventive controls.) If you need more SCPs on an account, you can nest OUs. For example, you can attach a maximum of 25 SCPs when you nest 5 OUs.

Automating controls

AWS Control Tower supports operational concurrency for all controls. That is, you can activate or deactivate multiple preventive and detective controls without having to wait for control operations to complete.

You can automatically activate and deactivate controls by using any of the following with the AWS Control Tower API:

- AWS CloudFormation
- AWS Command Line Interface (AWS CLI)
- Language-specific AWS SDKs

For more information about automating controls, see About controls in AWS Control Tower in the AWS Control Tower documentation. The following sections discuss mandatory controls, optional controls, and custom controls in more detail.

Mandatory controls

Mandatory controls are enforced by AWS Control Tower to protect AWS Control Tower managed resources. You can't deactivate mandatory controls.

Documenting mandatory controls for your organization

In your landing zone design document, you can document the mandatory controls that AWS Control Tower enforces by using the following table format. You can extend this table with optional controls and custom controls, as discussed later in this section.

> **Note**
>
> AWS Control Tower controls are continuously updated. For the most up-to-date and complete list of controls, see Mandatory controls in the AWS Control Tower documentation.

Control	Guidance level	Behavior	Default OU	Purpose
Disallow Changes to Encryption Configuration for AWS Control Tower Created Amazon S3 Buckets in Log Archive	Mandatory	Preventive	Security OU	Protects the encryption configuration for buckets deployed by AWS Control Tower in the Log Archive account so that encryption cannot be turned off for sensitive logs.
Disallow Changes to Logging Configuration for AWS Control	Mandatory	Preventive	Security OU	Protects the logging configuration for buckets deployed by

Control	Guidance level	Behavior	Default OU	Purpose
Tower Created Amazon S3 Buckets in Log Archive				AWS Control Tower in the Log Archive account so that only AWS Control Tower can make changes to these configurations.
Disallow Changes to Bucket Policy for AWS Control Tower Created Amazon S3 Buckets in Log Archive	Mandatory	Preventive	Security OU	Protects the bucket policies for buckets deployed by AWS Control Tower in the Log Archive account. This helps ensure that only AWS Control Tower can edit the permissions for the centralized logs, and that sensitive logs are secured.

Control	Guidance level	Behavior	Default OU	Purpose
Disallow Changes to Lifecycle Configuration for AWS Control Tower Created Amazon S3 Buckets in Log Archive	Mandatory	Preventive	Security OU	Protects the lifecycle configuration for buckets deployed by AWS Control Tower in the Log Archive account so that logs are stored for the required amount of time.
Disallow Changes to Amazon CloudWatch Logs Log Groups set up by AWS Control Tower	Mandatory	Preventive	All OUs	Protects the retention policy for the CloudWatch logs set up by AWS Control Tower in the Log Archive account so that only AWS Control Tower can make changes and logs are secured.

Control	Guidance level	Behavior	Default OU	Purpose
Disallow Deletion of AWS Config Aggregation Authorizations Created by AWS Control Tower	Mandatory	Preventive	All OUs	Protects the AWS Config aggregation authorizations set up by AWS Control Tower in the Audit account. This helps ensure that only AWS Control Tower can modify or disable account authorizations and that all authorization changes can be logged.
Disallow Deletion of Log Archive	Mandatory	Preventive	Security OU	Prevents deletion of the S3 buckets created by AWS Control Tower in the Log Archive account. This helps ensure that no one can remove the central log buckets.

Control	Guidance level	Behavior	Default OU	Purpose
Detect Public Read Access Setting for Log Archive	Mandatory	Detective	Security OU	Detects changes to read access permissions to the bucket deployed by AWS Control Tower in the Log Archive account. Such changes could risk exposing the central logs to the public.
Detect Public Write Access Setting for Log Archive	Mandatory	Detective	Security OU	Detects changes to write access permissions to the bucket deployed by AWS Control Tower. Such changes could risk exposing the central logs to the public.

Control	Guidance level	Behavior	Default OU	Purpose
Disallow Configuration Changes to CloudTrail	Mandatory	Preventive	All OUs	Protects the configuration of the organization trail deployed by AWS Control Tower. This helps ensure that only AWS Control Tower can modify the trail.
Integrate CloudTrail Events with Amazon CloudWatch Logs	Mandatory	Preventive	All OUs	Protects the CloudTrail event selectors of the organization trail deployed by AWS Control Tower.
Enable CloudTrail in All Available Regions	Mandatory	Preventive	All OUs	Protects the configuration of the organization trail deployed by AWS Control Tower in all enabled AWS Regions. This helps ensure that CloudTrail always collects logs in all enabled Regions.

Control	Guidance level	Behavior	Default OU	Purpose
Enable Integrity Validation for CloudTrail Log File	Mandatory	Preventive	All OUs	Protects the integrity of CloudTrail log files in the organization trail deployed by AWS Control Tower. Enabling integrity validation helps ensure that the digest file created for the logs can always prove that logs have not been modified.
Disallow Changes to Amazon CloudWatch Set Up by AWS Control Tower	Mandatory	Preventive	All OUs	Protects the CloudWatch logs set up by AWS Control Tower from modificat ion or removal so that AWS Control Tower log configura tions aren't modified.

Control	Guidance level	Behavior	Default OU	Purpose
Disallow Changes to Tags Created by AWS Control Tower for AWS Config Resources	Mandatory	Preventive	All OUs	Prevents changes to the tags that AWS Control Tower created when you set up the landing zone. This helps secure the AWS Control Tower functiona lity that is dependent on those tags.
Disallow Configuration Changes to AWS Config	Mandatory	Preventive	All OUs	Protects the AWS Config configuration set up by AWS Control Tower so that AWS Config recording cannot be modified or stopped.
Enable AWS Config in All Available Regions	Mandatory	Preventive	All OUs	Protects the AWS Config configuration set up by AWS Control Tower so that AWS Config recording cannot be modified or stopped in any AWS Region.

Control	Guidance level	Behavior	Default OU	Purpose
Disallow Changes to AWS Config Rules Set Up by AWS Control Tower	Mandatory	Preventive	All OUs	Protects the AWS Config Rules that are set up by AWS Control Tower to prevent them from being modified or removed. This helps ensure that the controls that are specific to AWS Control Tower are managed by AWS Control Tower only.
Disallow Changes to AWS IAM Roles Set Up by AWS Control Tower and AWS CloudFormation	Mandatory	Preventive	All OUs	Prevents changes to the IAM roles that AWS Control Tower created when you set up the landing zone so that the landing zone is secured.

Control	Guidance level	Behavior	Default OU	Purpose
Disallow Changes to AWS Lambda Functions Set Up by AWS Control Tower	Mandatory	Preventive	All OUs	Prevents changes to the AWS Lambda functions that are set up by AWS Control Tower so that the landing zone is secured.
Disallow Changes to Amazon SNS Set Up by AWS Control Tower	Mandatory	Preventive	All OUs	Prevents changes to the Amazon SNS topics that are set up by AWS Control Tower so that the landing zone is secured.
Disallow Changes to Amazon SNS Subscriptions Set Up by AWS Control Tower	Mandatory	Preventive	All OUs	Prevents changes to the Amazon SNS subscriptions that are set up by AWS Control Tower so that the integrity of Amazon SNS subscription settings for your landing zone are secured.

Control	Guidance level	Behavior	Default OU	Purpose
Detect whether shared accounts under the Security organizational unit have AWS CloudTrail or CloudTrail Lake enabled	Mandatory	Detective	Security OU	Detects whether AWS CloudTrail and AWS CloudTrail Lake are disabled in the accounts under the security OU.

Optional controls

You can enable optional controls on OUs in the organization if you choose. These controls are categorized as *strongly recommended* or *elective* controls. Strongly recommended controls are based on best practices for well-architected, multi-account environments. Elective controls prevent or track attempts to perform commonly restricted actions in an AWS enterprise environment. Unlike mandatory controls, strongly recommended and elective controls aren't activated by default—you can activate and deactivate them according to your requirements.

Security and compliance requirements

Make sure that you customize and adapt your control configurations and choices according to your landing zone requirements. The security requirements of your organization determine which controls to use and which OUs to enable them on. Before you select optional controls, you should consider your organization's specific goals, requirements, and compliance needs. Perform a comprehensive risk assessment to identify the specific risks and vulnerabilities that your organization faces in its AWS environment, and gather your security and compliance requirements. After you list your requirements clearly, you can start selecting the optional controls.

Guidelines

Strongly recommended controls are rooted in industry best practices for setting up a secure landing zone. Therefore, unless you have specific requirements that prevent their implementation, we recommend that you enable these controls across all OUs where the associated resources are provisioned.

Elective controls encompass industry-specific best practices and are tailored to address the unique security and compliance requirements of certain industries. We recommend that you research the best practices for your industry and adapt the relevant elective controls accordingly. The controls are designed to strengthen the security and compliance of your AWS environment, and adhering to them helps you align with recognized security standards.

However, some OUs might have unique circumstances that warrant exceptions. For example, consider enabling controls related to Amazon Elastic Block Store (Amazon EBS) volume encryption in OUs, such as workload OUs, where sensitive data is expected. Conversely, in a sandbox OU where experimentation is encouraged and no sensitive data is involved, you might have the flexibility to skip certain controls. The key is to balance robust security, compliance, and operational flexibility. Always aim to apply controls where they provide the most value while respecting the specific needs of each OU.

Documenting optional controls for your organization

You can use a table similar to the following in your design document to mark which optional controls should be enabled on which OUs. You can extend this table with information about the mandatory and custom controls you're using in your organization.

This table includes both strongly recommended and elective controls. The AWS Security Hub standard controls, data residency controls, and proactive controls are additional optional controls that you can append to the table. These are described later in this section.

The following table shows example configurations and OUs that you should adjust for your specific security and compliance requirements.

> **Note**
>
> AWS Control Tower controls are continuously updated. For the most up-to-date and complete list, see Optional controls in the AWS Control Tower documentation.

Control	Guidance level	Behavior	Security OU	Infrastructure OU	Suspended OU	Workloads OU	Deployments OU	Sandbox OU	Purpose
Disallow Creation of Access Keys for the Root User	Strongly recommended	Preventive	Yes	Yes	Yes	Yes	Yes	Yes	Reduces the risk of unauthorized access to the sensitive root user.
Disallow Actions as a Root User	Strongly recommended	Preventive	Yes	Yes	Yes	Yes	Yes	Yes	Reduces the impact of unauthorized access to the sensitive root user.
Detect Whether Encryption is Enabled for Amazon EBS Volumes Attached	Strongly recommended	Detective	Yes	Yes	Yes	Yes	Yes	No	Ensures that encryption is enabled to strengthen data security, maintain

Control	Guidance level	Behavior	Security OU	Infrastructure OU	Suspended OU	Workloads OU	Deployments OU	Sandbox OU	Purpose
to Amazon EC2 Instances									compliance, mitigate risks, or align with security best practices.
Detect Whether Unrestricted Incoming TCP Traffic is Allowed	Strongly recommended	Detective	Yes	Yes	Yes	Yes	Yes	No	Helps reduce the network attack surface for TCP traffic.
Detect Whether Unrestricted Internet Connection Through SSH is Allowed	Strongly recommended	Detective	Yes	Yes	Yes	Yes	Yes	No	Helps reduce the network attack surface for SSH traffic.

Control	Guidance level	Behavior	Security OU	Infrastructure OU	Suspended OU	Workloads OU	Deployments OU	Sandbox OU	Purpose
Detect Whether MFA for the Root User is Enabled	Strongly recommended	Detective	Yes	Yes	Yes	Yes	Yes	Yes	Helps reduce the risk of unauthorized access to the sensitive root user through multi-factor authentication.

Control	Guidance level	Behavior	Security OU	Infrastructure OU	Suspended OU	Workloads OU	Deployments OU	Sandbox OU	Purpose
Detect Whether Public Read Access to Amazon S3 Buckets is Allowed	Strongly recommended	Detective	Yes	Yes	Yes	Yes	Yes	No	Mitigates the risk of unauthorized read access to sensitive data by identifying S3 buckets that might be publicly accessible.

Control	Guidance level	Behavior	Security OU	Infrastructure OU	Suspended OU	Workloads OU	Deployments OU	Sandbox OU	Purpose
Detect Whether Public Write Access to Amazon S3 Buckets is Allowed	Strongly recommended	Detective	Yes	Yes	Yes	Yes	Yes	No	Mitigates the risk of unauthorized write access to sensitive data by identifying S3 buckets that might be publicly accessible.

Control	Guidance level	Behavior	Security OU	Infrastructure OU	Suspended OU	Workloads OU	Deployments OU	Sandbox OU	Purpose
Detect Whether Amazon EBS Volumes are Attached to Amazon EC2 Instances	Strongly recommended	Detective	Yes	Yes	Yes	Yes	Yes	No	Detects whether an Amazon EBS volume device persists independently from an Amazon EC2 instance.
Detect Whether Amazon EBS Optimization is Enabled for Amazon EC2 Instances	Strongly recommended	Detective	Yes	Yes	Yes	Yes	Yes	No	Detects EC2 instances where performance and cost can be improved by using Amazon EBS optimization.

Control	Guidance level	Behavior	Security OU	Infrastructure OU	Suspended OU	Workloads OU	Deployments OU	Sandbox OU	Purpose
Detect Whether Public Access to Amazon RDS Database Instances is Enabled	Strongly recommended	Detective	Yes	Yes	Yes	Yes	Yes	No	Detects publicly accessible Amazon Relational Database Service (Amazon RDS) database instances to secure sensitive data.
Detect Whether Public Access to Amazon RDS Database Snapshots is Enabled	Strongly recommended	Detective	Yes	Yes	Yes	Yes	Yes	Yes	Detects publicly accessible Amazon RDS database snapshots to secure sensitive data.

Control	Guidance level	Behavior	Security OU	Infrastructure OU	Suspended OU	Workloads OU	Deployments OU	Sandbox OU	Purpose
Detect Whether Storage Encryption is Enabled for Amazon RDS Database Instances	Strongly recommended	Detective	Yes	Yes	Yes	Yes	Yes	No	Identifies unencrypted Amazon RDS instances to mitigate risk of sensitive data exposure.
Detect whether an account has AWS CloudTrail or CloudTrail Lake enabled	Strongly recommended	Detective	Yes	Yes	Yes	Yes	Yes	Yes	Ensures that proper monitoring is enabled by using CloudTrail.

Control	Guidance level	Behavior	Security OU	Infrastructure OU	Suspended OU	Workloads OU	Deployments OU	Sandbox OU	Purpose
Disallow Changes to Replication Configuration for Amazon S3 Buckets	Elective	Preventive	Yes	Yes	Yes	Yes	Yes	No	Prevents unauthorized alterations to replication configurations to ensure consistent data replication and adherence to regulatory requirements.

Control	Guidance level	Behavior	Security OU	Infrastructure OU	Suspended OU	Workloads OU	Deployments OU	Sandbox OU	Purpose
Disallow Delete Actions on Amazon S3 Buckets Without MFA	Elective	Preventive	Yes	Yes	Yes	Yes	Yes	No	Prevents accidental or malicious deletion of S3 buckets by requiring multi-factor authentication.
Detect Whether MFA is Enabled for AWS IAM Users	Elective	Detective	Yes	Yes	Yes	Yes	Yes	No	Identifies IAM users that don't have multi-factor authentication enabled, to mitigate the risk of unauthorized access.

Control	Guidance level	Behavior	Security OU	Infrastructure OU	Suspended OU	Workloads OU	Deployments OU	Sandbox OU	Purpose
Detect Whether MFA is Enabled for AWS IAM Users of the AWS Console	Elective	Detective	Yes	Yes	Yes	Yes	Yes	No	Identifies IAM users in the AWS Management Console that don't have multi-factor authentication enabled, to mitigate the risk of unauthorized access.

Control	Guidance level	Behavior	Security OU	Infrastructure OU	Suspended OU	Workloads OU	Deployments OU	Sandbox OU	Purpose
Detect Whether Versioning for Amazon S3 Buckets is Enabled	Elective	Detective	Yes	Yes	Yes	Yes	Yes	No	Identifies S3 buckets where versioning isn't enabled, to mitigate the risk of accidental deletion or modification of data.
Disallow Changes to Encryption Configuration for Amazon S3 Buckets	Elective	Preventive	Yes	Yes	Yes	Yes	Yes	No	Prevents changes to encryption configuration of S3 buckets to protect sensitive data.

Control	Guidance level	Behavior	Security OU	Infrastructure OU	Suspended OU	Workloads OU	Deployments OU	Sandbox OU	Purpose
Disallow Changes to Logging Configuration for Amazon S3 Buckets	Elective	Preventive	Yes	Yes	Yes	Yes	Yes	No	Prevents changes to logging configuration for S3 buckets to ensure consistent and reliable audit logging.
Disallow Changes to Bucket Policy for Amazon S3 Buckets	Elective	Preventive	Yes	Yes	Yes	Yes	Yes	No	Prevents changes to bucket policies for S3 buckets to maintain proper access controls.

Control	Guidance level	Behavior	Security OU	Infrastructure OU	Suspended OU	Workloads OU	Deployments OU	Sandbox OU	Purpose
Disallow Changes to Lifecycle Configuration for Amazon S3 Buckets	Elective	Preventive	Yes	Yes	Yes	Yes	Yes	No	Prevents changes to lifecycle configurations for S3 buckets to help maintain data management consistency and compliance.

Control	Guidance level	Behavior	Security OU	Infrastructure OU	Suspended OU	Workloads OU	Deployments OU	Sandbox OU	Purpose
Disallow management of resource types, modules, and hooks within the AWS CloudFormation registry	Elective	Preventive **(Note:** You must enable this control when you activate proactive controls in your environment.)	Yes	Yes	Yes	Yes	Yes	Yes	Prevents unintended management of resource type, modules, and hooks to help ensure the stability and security of infrastructure deployments.

AWS Security Hub controls

AWS Control Tower is integrated with AWS Security Hub through a Security Hub standard. This integration provides additional controls that help you streamline security and compliance management in your AWS environment.

You can combine more than 230 detective controls from Security Hub with AWS Control Tower controls to help cover your security and compliance requirements. You can add your selected controls to the table that you set up in the previous section.

Note

To start using Security Hub controls in AWS Control Tower, go to the AWS Control Tower controls library and enable the desired Security Hub control. AWS Control Tower takes care of the activation process and creates a new standard named **Service-Managed Standard: AWS Control Tower** in Security Hub. This standard provides visibility into activated controls and their evaluations, which simplifies monitoring and compliance efforts. For more information, see Security Hub standard in the AWS Control Tower documentation.

Data residency controls

Data residency controls enforce data residency requirements in your organization. These elective controls are included in AWS Control Tower to help ensure that your data is stored and processed in compliance with your regulations and policies. You should consider using data residency controls in scenarios such as the following:

- **Regulatory compliance:** You want to ensure that data is stored and processed in the designated geographic regions to meet regulatory requirements such as General Data Protection Regulation (GDPR), Health Insurance Portability and Accountability Act (HIPAA), or industry-specific regulations.
- **International operations:** You want to segment your AWS workloads based on their geographic locations and ensure that data remains within the desired region.
- **Risk mitigations:** You want to mitigate the risk of data exposure from accidental or unauthorized data transfers across regions, to reduce the risk of data leakage or non-compliance.
- **Data sovereignty:** You run workloads in countries that have laws that require data to remain within the country's borders.
- **Data classification:** You want to classify data based on its sensitivity or regulatory requirements, and then apply specific policies to each data classification.

It is essential to thoroughly understand your organization's data residency requirements and the relevant regulations before implementing data residency controls in AWS Control Tower.

Documenting data residency controls for your organization

When you design your data residency controls, you can use the optional controls table provided previously in this section and append the data residency controls that you have selected to meet

your requirements. The following table lists the existing controls and examples of when to use them.

> **Note**
>
> AWS Control Tower controls are continuously updated. For the most up-to-date and complete list of controls, see Controls that enhance data residency protection in the AWS Control Tower documentation.

Control	Guidance level	Behavior	Default OU	Purpose
Deny access to AWS based on the requested AWS Region	Elective	Preventive	All OUs, if enabled in AWS Control Tower landing zone settings.	(This control is frequently referred to as the Region deny control.) Ensures that AWS resources are provisioned only in approved AWS Regions, aligning with data residency and compliance requirements.
Disallow internet access for an Amazon VPC instance managed by a customer	Elective	Preventive	—	Prevents internet access in VPCs to reduce the risk of unauthori zed access or data exposure to the public when there are

Control	Guidance level	Behavior	Default OU	Purpose
				data residency and privacy requirements.
Disallow AWS Virtual Private Network (AWS VPN) connections	Elective	Preventive	—	Restricts VPN connections to guard against unauthorized access, data exfiltration, or bypassing security controls.
Disallow cross-region networking for Amazon EC2, Amazon CloudFront, and AWS Global Accelerator	Elective	Preventive	—	Prevents cross-Region networking to maintain data residency and help ensure that data remains within approved Regions. Public access could inadvertently lead to data being distributed outside these boundaries.

Control	Guidance level	Behavior	Default OU	Purpose
Detect whether public IP addresses for Amazon EC2 autoscaling are enabled through launch configurations	Elective	Detective	—	Monitors and controls the exposure of instances to the public internet. This helps reduce the attack surface and risk of unauthorized access that might compromise data residency and security.
Detect whether replication instances for AWS Database Migration Service are public	Elective	Detective	—	Ensures that replication instances aren't publicly accessible, which helps protect sensitive data from unauthorized access and data residency violations.

Control	Guidance level	Behavior	Default OU	Purpose
Detect whether Amazon EBS snapshots are restorable by all AWS accounts	Elective	Detective	—	Limits access to EBS snapshots to help prevent unauthorized access, data breaches, and potential non-compliance with data residency regulations.
Detect whether any Amazon EC2 instance has an associated public IPv4 address	Elective	Detective	—	Helps identify and mitigate security risks associated with instances that have public IP addresses. These instances might be more vulnerable to attacks.
Detect whether Amazon S3 settings to block public access are set as true for the account	Elective	Detective	—	Enforces strict access controls on Amazon S3 buckets to prevent unauthorized public access to sensitive data, to align with data residency and privacy needs.

Control	Guidance level	Behavior	Default OU	Purpose
Detects whether an Amazon EKS endpoint is blocked from public access	Elective	Detective	—	Ensures that Amazon Elastic Kubernetes Service (Amazon EKS) cluster endpoints aren't accessible from the public internet. This helps prevent unauthorized sharing of sensitive data that might compromise data residency requirements.

Control	Guidance level	Behavior	Default OU	Purpose
Detect whether an Amazon OpenSearch Service domain is in Amazon VPC	Elective	Detective	—	Ensures that Amazon OpenSearch Service domain endpoints aren't public. Deploying these domains within VPCs improves data security by preventing public access and maintaining data residency within trusted network boundaries.
Detect whether any Amazon EMR cluster master nodes have public IP addresses	Elective	Detective	—	Reduces security risks of compromising data residency requirements by ensuring that Amazon EMR cluster master nodes don't have publicly accessible IP addresses.

Control	Guidance level	Behavior	Default OU	Purpose
Detect whether the AWS Lambda function policy attached to the Lambda resource blocks public access	Elective	Detective	—	Controls access to AWS Lambda functions and prevents unauthorized public invocation or exposure of sensitive functions.
Detect whether public routes exist in the route table for an Internet Gateway (IGW)	Elective	Detective	—	Helps maintain network security by ensuring that public routes through an internet gateway are configured only where necessary.
Detect whether Amazon Redshift clusters are blocked from public access	Elective	Detective	—	Ensures that Amazon Redshift clusters aren't publicly accessible. This helps protect clusters from unauthorized access that could compromise data residency.

Control	Guidance level	Behavior	Default OU	Purpose
Detect whether an Amazon SageMaker AI notebook instance allows direct internet access	Elective	Detective	—	Helps prevents direct internet access to SageMaker AI notebook instances to align with data residency and security requirements, and to reduce exposure to potential threats.
Detect whether any Amazon VPC subnets are assigned a public IP address	Elective	Detective	—	Helps maintain network isolation to reduce the risk of unauthorized data exposure and data residency violations.

Control	Guidance level	Behavior	Default OU	Purpose
Detect whether AWS Systems Manager documents owned by the account are public	Elective	Detective	—	Helps ensure that Systems Manager documents aren't publicly accessible. This helps protect sensitive data and maintain data residency and security.

Proactive controls

Proactive controls are optional controls that are implemented with AWS CloudFormation Hooks. This mechanism enables you to run custom logic during the deployment of CloudFormation stacks to monitor and validate the configuration settings and resources that are defined in the CloudFormation templates. If proactive controls detect any deviations or non-compliance issues, they can take immediate action, such as halting the deployment, sending notifications, or initiating remediation processes, to help mitigate potential risks and maintain the desired security posture.

Proactive controls in AWS Control Tower help you identify and address issues before they become vulnerabilities or compliance violations, and ensure a robust and well-governed AWS environment. These controls are designed to complement the existing guardrails and controls within AWS Control Tower. They can provide an additional layer of security and compliance assurance, especially in scenarios where early prevention and continuous monitoring are essential. However, the specific proactive controls you choose to implement should align with your organization's goals, risk profile, and compliance needs. If your organization has specific security requirements that go beyond the default AWS Control Tower controls, you can customize proactive controls to meet these needs.

These controls are categorized by service and listed in the Proactive controls section of the AWS Control Tower documentation. You can choose from a large selection of controls and add them to your selected controls table.

> **Note**
>
> AWS CloudFormation Hooks isn't supported in all AWS Regions where AWS Control Tower is available. Therefore, when you deploy a proactive control, it might not operate in all AWS Regions that you govern with AWS Control Tower.

Custom controls

After you have conducted your risk assessment, identified your security and compliance requirements, and selected the AWS Control Tower controls to guardrail these requirements, there might be some requirements that still aren't addressed. You can implement custom service control policies (SCPs), AWS Config Rules, and AWS CloudFormation Hooks to cover these requirements. However, these controls aren't implemented as AWS Control Tower controls—they're implemented outside AWS Control Tower.

The following table provides examples of custom controls that you can append to your controls table.

Control	Guidance level	Behavior	Security OU	Infrastructure OU	Suspended OU	Workloads OU	Deployments OU	Sandbox OU	Purpose
Protect Amazon CloudWatch	Custom SCP	Proactive	Yes	Yes	Yes	Yes	Yes	No	Deny `cloudwatch:DeleteAlarms` ,`cloudwatch:DeleteDashboards` , `cloudwatch:DisableAlarmActions` , `cloudwatch:PutDashboard` ,

Control	Guidance level	Behavior	Security OU	Infrastructure OU	Suspended OU	Workloads OU	Deployments OU	Sandbox OU	Purpose
									`cloudwatch:PutMetricAlarm`, `cloudwatch:SetAlarmState`
Enforce encryption for Amazon Simple Storage Service (Amazon S3) buckets	Custom SCP	Proactive	Yes	Yes	Yes	Yes	Yes	No	Deny `s3:PutObject` on the condition that encryption is `false`
AWS Identity and Access Management (IAM) user creation	Custom SCP	Proactive	Yes	Yes	Yes	Yes	Yes	Yes	Deny `iam:CreateUser`

Control	Guidance level	Behavior	Security OU	Infrastructure OU	Suspended OU	Workloads OU	Deployments OU	Sandbox OU	Purpose
Protect account and billing settings	Custom SCP	Proactive	Yes	Yes	Yes	Yes	Yes	Yes	Deny `aws-portal:ModifyAccount` , `aws-portal:ModifyBilling` , `aws-portal:ModifyPaymentMethods`

Networking integration

Most enterprises require connectivity between accounts in their AWS Control Tower–managed environment. This often extends to connecting corporate offices and on-premises data centers. AWS Virtual Private Network (AWS VPN) and AWS Direct Connect are used as network paths to provide that hybrid connectivity for workloads.

- AWS VPN establishes a secure and private tunnel from your network or device to the AWS Cloud over the internet. It allows you to securely connect your on-premises network or branch office site to your VPC.
- AWS Direct Connect makes it easy to establish a dedicated network connection from your on-premises environment to AWS. It provides a more consistent network experience than internet-based connections.
- AWS Transit Gateway connects VPCs and on-premises networks through a central hub and enables various routing scenarios. It controls how traffic is routed among the connected networks.

The easiest way to get started with hybrid connectivity is to establish site-to-site VPN over the internet. This extends your data center or branch office to the cloud by using IPsec tunnels. You can configure routing by using Border Gateway Protocol (BGP) or configure static routes. Each AWS Site-to-Site VPN connection consists of two VPN tunnel endpoints for redundancy. Each tunnel terminates in a different Availability Zone within the AWS global network, for high availability.

AWS Site-to-Site VPN supports terminating IPsec tunnels on both virtual private gateways and AWS Transit Gateway at the AWS end. When you terminate a VPN on a virtual private gateway, you can access the VPC that the gateway is attached to. However, if you use Transit Gateway, you gain connectivity to thousands of VPCs over a pair of VPN tunnels. Additionally, Transit Gateway supports equal-cost multipath (ECMP) routing, which enables you to load-balance traffic across multiple VPN tunnels for high availability and bandwidth aggregation.

In summary, terminating a VPN at a transit gateway is a default starting point for hybrid architectures, because it provides more flexibility in the number of VPCs you can connect to, and added functionality such as ECMP.

The following diagram shows how you can connect an on-premises environment to your VPCs on AWS by using AWS Site-to-Site VPN.

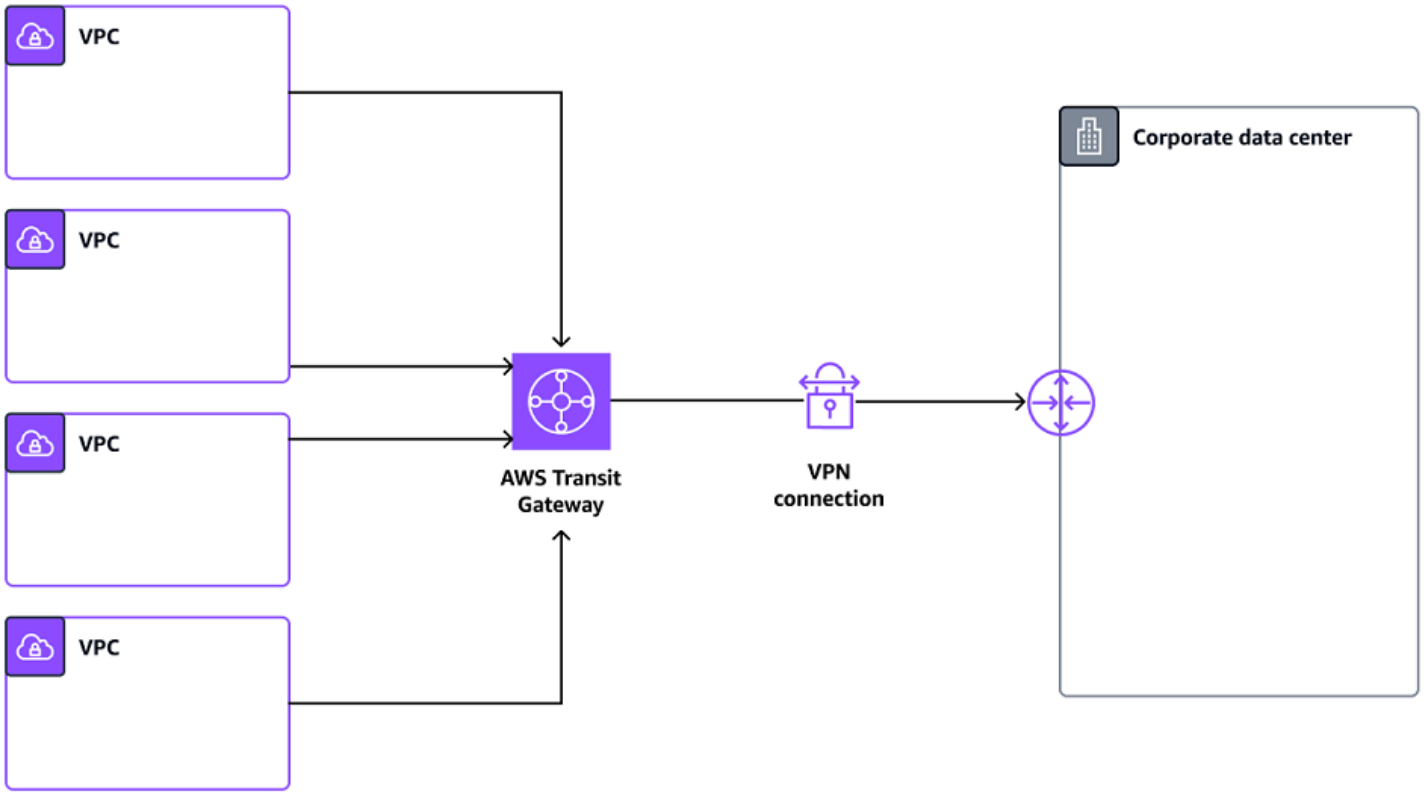

For end-to-end network performance, you can use AWS Direct Connect to enable consistent, low-latency, high-bandwidth, dedicated fiber connectivity between your on-premises data centers and AWS. AWS Direct Connect provides dedicated connections at bandwidths of 1 Gbps, 10 Gbps, 100 Gbps, and 400 Gbps. Hosted connections provided by AWS Direct Connect Partners use pre-established network links and are available from 50 Mbps up to 25 Gbps.

AWS Direct Connect provides three types of virtual interfaces (VIFs):

- Public VIFs provide global connectivity to public AWS resources, including AWS public service endpoints, public Amazon EC2 IP addresses, and public Elastic Load Balancing addresses.
- Private VIFs provide connectivity to the private IP range of your VPC.
- Transit VIFs enable connectivity to transit gateways.

The following sections provide examples of these connectivity options.

Sections:

- AWS Direct Connect with private VIF over virtual private gateway
- AWS Direct Connect with AWS Transit Gateway over transit VIF
- Inter-VPC connectivity through AWS Transit Gateway

- AWS Direct Connect SiteLink
- AWS Cloud WAN

AWS Direct Connect with private VIF over virtual private gateway

The following diagram shows how you can connect VPCs and on-premises environments through a virtual private gateway over a private VIF by using AWS Direct Connect.

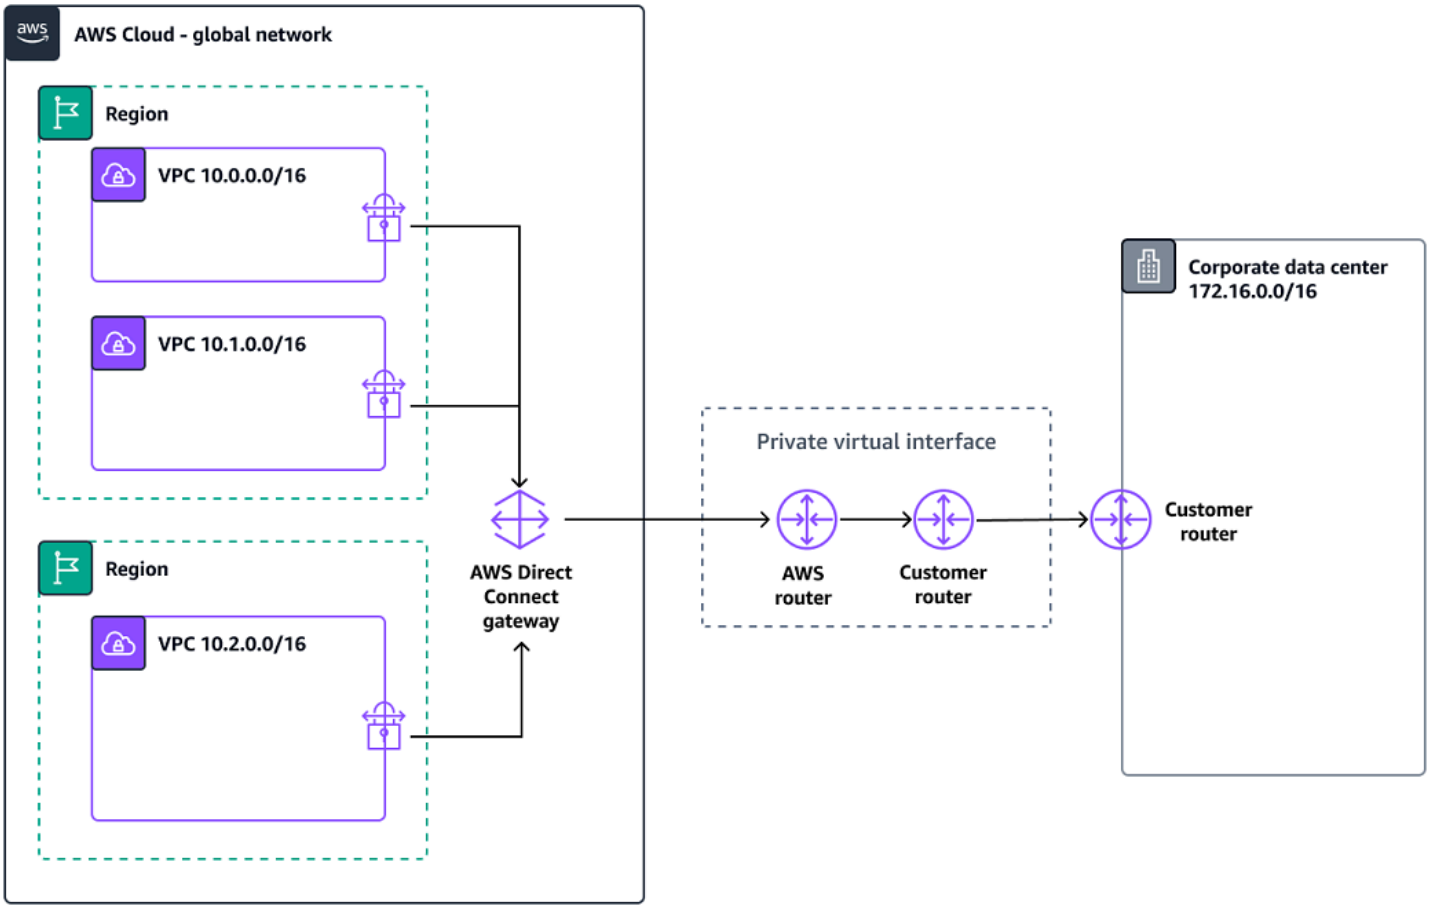

Most large enterprise customers deploy resources within a large number of VPCs across multiple AWS Regions and require connectivity from data centers that are spread across geographies. By using an AWS Direct Connect gateway, which is a global construct, you can use existing AWS Direct Connect connections to connect to resources in VPCs across AWS Regions. You can associate up to 10 virtual private gateways (each attached to a VPC) in different AWS Regions, directly to an AWS Direct Connect gateway. Alternatively, you can create a transit VIF and attach a total of six transit

gateways (each attached to thousands of VPCs) to an AWS Direct Connect gateway across AWS Regions.

AWS Direct Connect with AWS Transit Gateway over transit VIF

The following diagram shows how you can connect VPCs from multiple AWS Regions to an on-premises environment by using AWS Direct Connect and AWS Transit Gateway over a transit VIF.

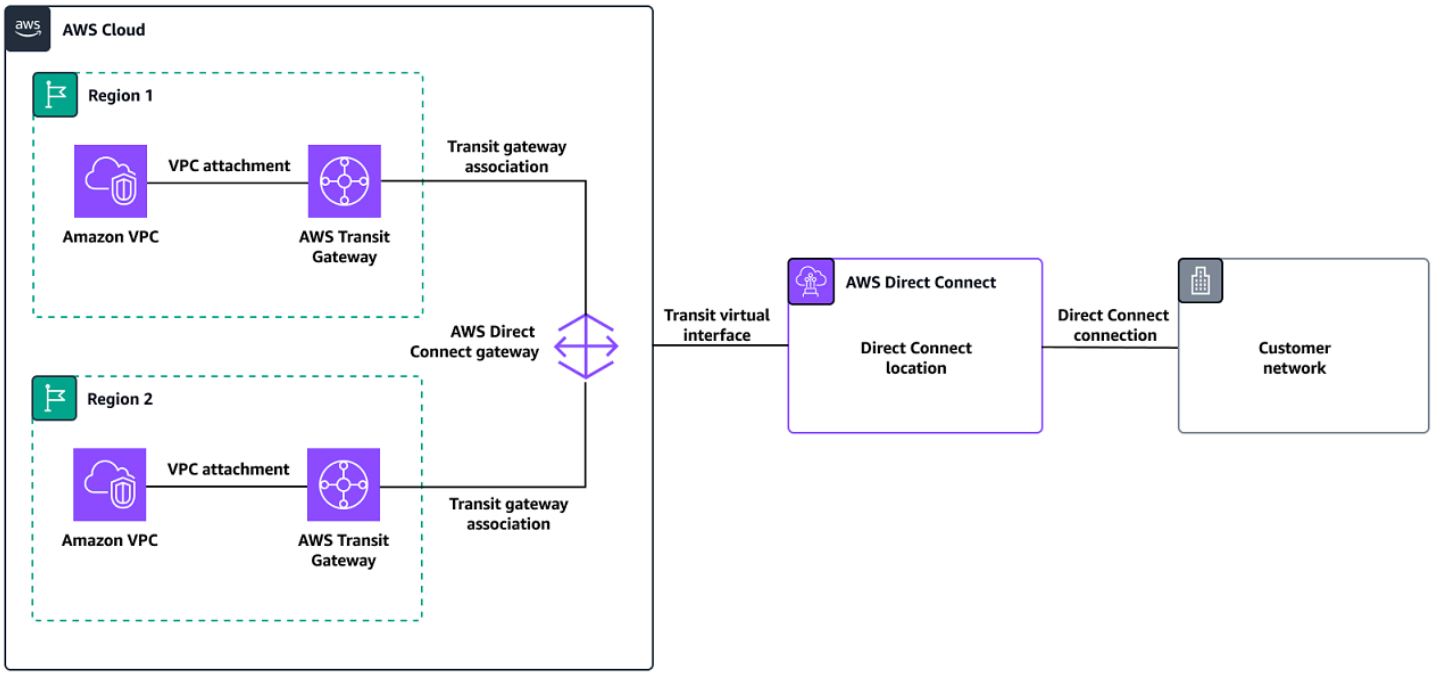

The transit gateway routes traffic through the centralized AWS Direct Connect gateway for all AWS Regions. A transit VIF attachment to the AWS Direct Connect gateway enables your network to connect up to six Regional, centralized transit gateways over a private, dedicated connection.

Inter-VPC connectivity through AWS Transit Gateway

The following diagram shows how you can interconnect VPCs through a transit gateway in the same AWS Region.

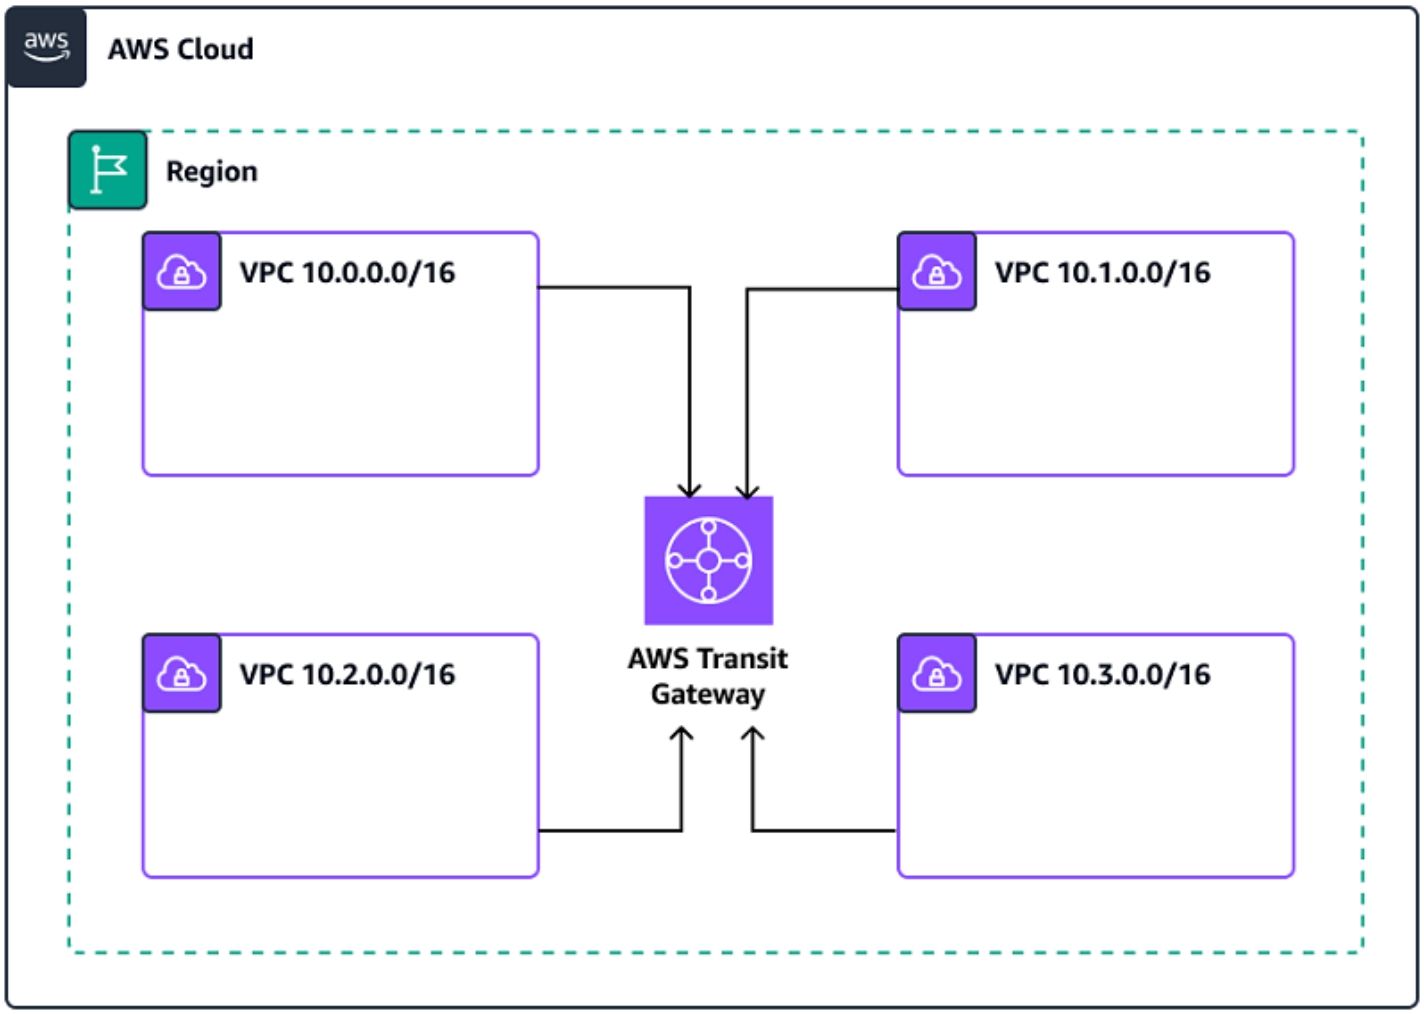

A transit gateway is a network transit hub that you can use to connect your VPCs and on-premises networks. As your organization grows, you can peer transit gateways from different AWS Regions to allow connectivity between them.

AWS Direct Connect SiteLink

If your requirements include sending or routing traffic directly between AWS Direct Connect locations, you can use AWS Direct Connect SiteLink. This feature of AWS Direct Connect sends traffic between different AWS Direct Connect locations over the shortest path on the AWS network without entering AWS Regions.

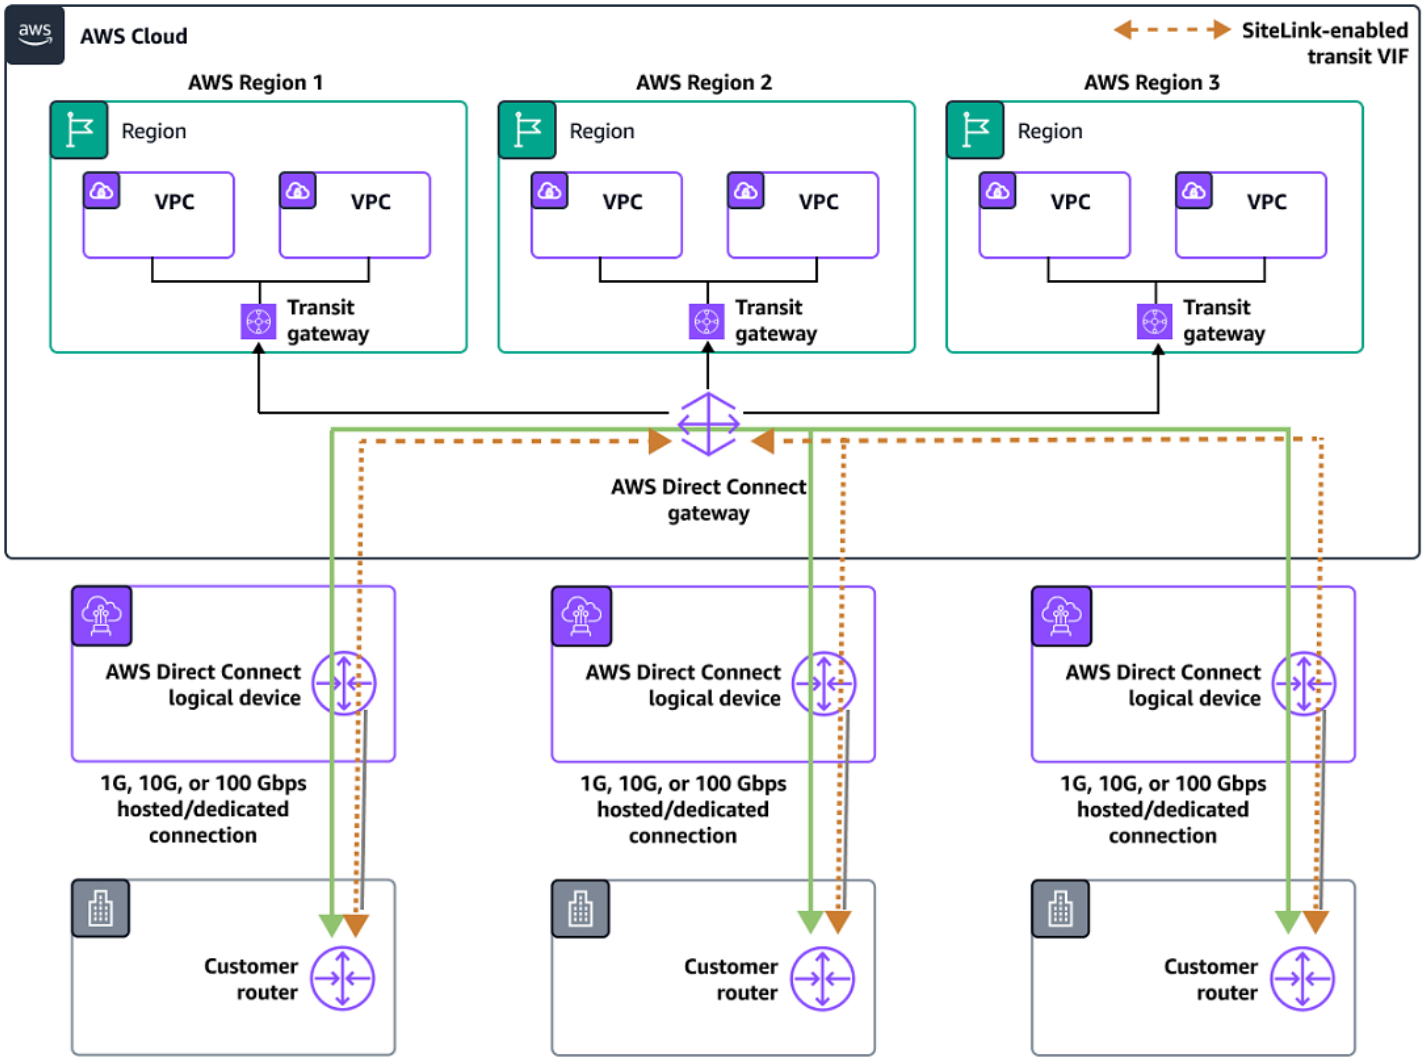

AWS Cloud WAN

Although you can create your own global network by interconnecting multiple transit gateways across Regions, you can also take advantage of AWS Cloud WAN. This service provides built-in automation, segmentation, and configuration management features that are designed specifically for building and operating global networks, based on your core network policy.

Both AWS Transit Gateway and AWS Cloud WAN allow centralized connectivity between VPCs and on-premises locations. Transit Gateway is a Regional network connectivity hub and is optimal if you operate in a few AWS Regions and want to manage your own peering and routing configuration. AWS Cloud WAN is optimal for users who want to define their global network through policy and have the service implement the underlying components automatically.

Authentication and authorization

Central cloud administrators and end users can use AWS IAM Identity Center to manage access to multiple AWS accounts and business applications. When you set up a landing zone, AWS Control Tower gives you two options for authentication:

- AWS managed account access with IAM Identity Center
- Self-managed AWS account access with IAM Identity Center or another method

If you are setting up a new landing zone and would like IAM Identity Center to set up and manage access to your accounts based on AWS best practices, choose the first option. If you have an existing landing zone and use IAM Identity Center or a third-party identity provider, choose the second option. In this case, you have to install AWS Control Tower in the same AWS Region as your existing IAM Identity Center deployment. If you're using an existing IAM Identity Center identity source, AWS Control Tower won't delete or modify your configuration. You will still be able to manage any further changes to IAM Identity Center configuration yourself.

For a new landing zone, if you choose the AWS Control Tower setup option, you can choose an IAM Identity Center directory, a SAML 2.0-compatible identity provider (IdP), or Active Directory as your identity source in IAM Identity Center. The identity source defines where you administer and authenticate identities. IAM provides these features:

- Active Directory users and user groups are synchronized between your identity source and IAM Identity Center.
- AWS permission sets are defined by job roles, such as infrastructure administrator or security operations.
- User groups from the identity source are mapped to the defined permissions.
- You must require multi-factor authentication (MFA) for all root, IAM, and IAM Identity Center users.

You can also use an external identity provider as your identity source to manage access to your AWS accounts, resources, and cloud applications. During SAML-based authentication, users and groups are synchronized from your external identity provider by using System for Cross-domain Identity Management (SCIM) in IAM Identity Center. Users can complete this federation by using the IAM Identity Center portal. The following diagram illustrates how identify federation works.

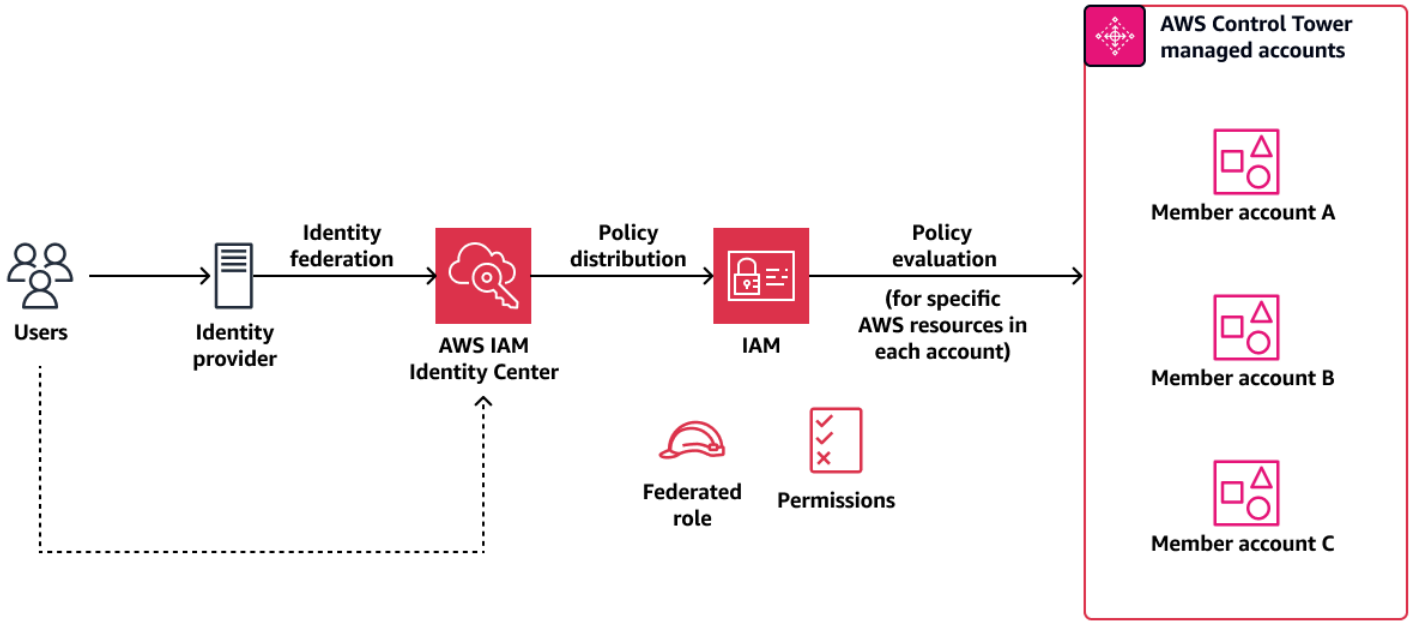

Direct access to AWS accounts must be limited only through the AWS account root user and break glass identities (IAM roles or users that can access the accounts if the IAM Identity Center federation is broken or you are accidentally locked out of the environment).

Break glass access

Break glass access refers to a quick means for a person who doesn't have access privileges to certain AWS accounts to gain access in exceptional circumstances, by using an approved process. The management account in AWS Organizations is used to provide break glass access to AWS accounts within the organization.

AWS discourages the use and creation of IAM users. However, break glass users are an exception. These users assume roles in the member accounts in your organization through trust policies. A break glass role that only the break glass users from the management account can assume is deployed to all the accounts in the organization. When you set up these roles in your organization, make sure that they can be used in emergency situations, such as the failure of the organization's identity provider, security incidents, or unavailability of key personnel, to provide temporary, elevated access beyond regular permissions to perform tasks such as updating guardrails, troubleshooting issues with automation tooling, or remediating security and operational issues that might occur. For more information, see Set up emergency access to the AWS Management Console in the IAM Identity Center documentation.

> **Warning**
>
> IAM users have long-term credentials, which present a security risk. To help mitigate this risk, we recommend that you provide these users with only the permissions they require to perform the task and that you remove these users when they are no longer needed.

Roles and responsibilities

Here's the approach for granting the appropriate level of access to new accounts in your landing zone:

- New groups are created in your IdP according to the required job function. For example, an `AWS-Management-BillingAdmin` group could modify billing information, modify payment methods, download invoices, and read from AWS Cost Explorer in the management account, but wouldn't be able to access other accounts.
- The groups created in the identity source are visible in IAM Identity Center after federation is complete.
- You can define new permission sets in IAM Identity Center. A permission set defines the level of access that users and groups have to an AWS account. They are stored in IAM Identity Center and can be provisioned to one or more AWS accounts. For example, you could create a `BillingAdmin` permission set for the `AWS-Management-BillingAdmin` group.

  > **Note**
  >
  > IAM Identity Center provides predefined permission sets such as `AWSReadOnlyAccess` and `AWSAdministratorAccess`.

- IAM Identity Center provides AWS managed policies for job functions through an IAM policy that provides the appropriate level of access to AWS services. You can attach these managed policies to permission sets in IAM Identity Center. For example, you can attach the `Billing` managed policy to the `BillingAdmin` permission set. You can also create custom policies, if required.
- In IAM Identity Center, you associate accounts with an identity source group and permission set. For example, you can associate the management account with the `AWS-Management-BillingAdmin` group and `BillingAdmin` permission set.

The following table lists the AWS managed policies for job functions that are available in IAM Identity Center. You can use these as a starting point for defining permission sets. For more information, see AWS managed policies for job functions in the IAM documentation.

AWS managed policy name	Description of job function
`AdministratorAccess`	Provides full access to AWS services and resources.
`Billing`	Grants permissions for billing and cost management. This includes viewing account usage and viewing or modifying budgets and payment methods.
`DataScientist`	Grants permissions to AWS data analytics services.
`DatabaseAdministrator`	Grants full access permissions to the AWS services and actions required to set up and configure AWS database services.
`NetworkAdministrator`	Grants full access permissions to the AWS services and actions required to set up and configure AWS network resources.
`PowerUserAccess`	Provides full access to AWS services and resources for application developers, but doesn't allow management of users and groups.
`SecurityAudit`	Grants read access to the security configura tion metadata. This is useful for software that audits the configuration of an AWS account.
`SupportUser`	Grants permissions to troubleshoot and resolve issues in an AWS account. This policy also enables the user to contact AWS Support.

AWS managed policy name	Description of job function
SystemAdministrator	Grants full access permissions to the resources required for application and development operations.
ViewOnlyAccess	Grants permissions to view resources and basic metadata across all AWS services.

The following table describes the additional permission sets that you can set up in IAM Identity Center, along with the associated accounts.

> **Important**
>
> Make sure that you choose and customize your permission sets according to your landing zone requirements.

Permission set created in IAM Identity Center	AWS managed policies for job functions	Active Directory group	AWS account associated with the Active Directory group
Administrator	AdministratorAccess	AWS-Management-Administrator	Management
		AWS-Audit-Administrator	Audit
		AWS-LogArchive-Administrator	Log Archive
		AWS-SharedServices-Administrator	Shared Services

Permission set created in IAM Identity Center	AWS managed policies for job functions	Active Directory group	AWS account associated with the Active Directory group
		AWS-Networking-Administrator	Networking
		AWS-Prod-Administrator	Production
		AWS-NonProd-Administrator	Non-production
BillingAdmin	Billing	AWS-Management-BillingAdmin	Management
SecurityAuditor	SecurityAudit	AWS-Global-SecurityAuditor	All accounts
ReadOnly	ViewOnlyAccess	AWS-Core-ReadOnly	Audit, Log Archive
		AWS-Infrastructure-ReadOnly	Shared Services, Networking
		AWS-Infrastructure-ReadOnly-NonProd	Non-production
NetworkPowerUser	NetworkAdministrator	AWS-Infrastructure-NetworkPowerUser	Networking

Permission set created in IAM Identity Center	AWS managed policies for job functions	Active Directory group	AWS account associated with the Active Directory group
Support	SupportUser	AWS-Global-SupportUser	Log Archive, Shared Services, Networking, Non-production

Centralized logging and monitoring

Organizations often create dedicated AWS accounts for centralized logging and monitoring purposes. These accounts are used to collect and store logs from various AWS accounts and services within the organization for long-term archival and auditing as well as monitoring the activity in all accounts for threats and vulnerabilities. In the security OU, AWS Control Tower implements a centralized log store (Log Archive) for logs and a centralized audit account (Audit) for auditor access and security tooling.

> **Note**
>
> You might have customized these default account and OU names while setting up your landing zone in AWS Control Tower.

Logging, monitoring, and alerting are important components of an AWS Control Tower landing zone. Some functionalities are automatically launched when you set up the landing zone, and you can add other functionalities later for a more comprehensive landing zone monitoring solution.

Topics

Logging

The Log Archive account serves as a centralized repository for aggregating logs of API activities (by using AWS CloudTrail) and resource configurations (by using AWS Config) across all accounts within the landing zone. Furthermore, you can centralize other logs from across your organization, such as Amazon CloudWatch, Amazon S3 access logs, and VPC Flow Logs, in this account. The Log Archive account seamlessly integrates with AWS Control Tower to automatically capture and record actions and events. This includes actions initiated from both the management account and member accounts. For comprehensive guidance, see Logging and monitoring in AWS Control Tower in the AWS Control Tower documentation.

Centralized logging in AWS Control Tower provides numerous benefits, including:

- Integration of security services to audit the logs and automate alerts and remediations
- Adherence to compliance and regulatory standards that require you to keep a record of all activities in your environment
- Centralized visibility into all activities across accounts to enable rapid troubleshooting and aid in forensic analysis during security incidents
- Support for growing log volumes and cost-effective, long-term storage solutions

> **Note**
>
> To further enhance your centralized logging solution, you can use AWS solutions such as Centralized Logging with OpenSearch, which provides capabilities to ingest, process, and visualize both application logs and AWS service logs.

The following table provides an overview of the logs that you can set up for your landing zone, as an example of a table that you can use in your landing zone design document. You can extend this table with additional log solutions according to your landing zone requirements. For more guidance about the security logs to include in the Log Archive account, see the AWS Security Reference Architecture.

Logging service	Description	Build approach	Location
AWS CloudTrail and AWS Config	AWS Config logs configuration activity in the resources it supports. CloudTrail logs API calls, console access, and logins. Logs from all accounts are aggregated in the Log Archive account.	Automatically enabled and set up by AWS Control Tower for all accounts in the landing zone.	S3 bucket in the Log Archive account.
Amazon CloudWatch	CloudWatch monitors resources and applications in the	We recommend that you set up CloudWatc	S3 bucket configuration details are

Logging service	Description	Build approach	Location
	environment in real time. CloudWatch collects and tracks metrics for resources and applications.	h for all required AWS resources.	provided with the workloads.
Amazon S3 access logs	Amazon S3 access logging provides detailed records for requests made to an S3 bucket. AWS Control Tower automatically sets up Amazon S3 access logging in the S3 bucket for CloudTrail and AWS Config. For information about Amazon S3 access logging, see Logging requests using server access logging in the Amazon S3 documentation.	Automatically enabled and set up by AWS Control Tower in the S3 bucket for CloudTrail and AWS Config.	S3 bucket in the Log Archive account.

Logging service	Description	Build approach	Location
Elastic Load Balancing (ELB) access logs	ELB access logs capture detailed information about requests sent to your load balancer. These logs can be collected in all member accounts that have load balancers and centralized in the Log Archive bucket. For more informati on about ELB access logging, see Access logs for your Network Load Balancer and Access logs for your Application Load Balancer in the ELB documentation.	We recommend that you set up access logs for all ELB resources.	S3 bucket in the Log Archive account.

Logging service	Description	Build approach	Location
VPC Flow Logs	VPC Flow Logs captures informati on about IP traffic going to and from network interfaces in the VPC. These logs are locally stored in each member account and can be used for troubleshooting and analysis. For more informati on about this feature, see VPC Flow Logs in the Amazon VPC documentation.	We recommend that you use an AWS CloudFormation script to enable VPC Flow Logs when you set up a VPC in each account.	Locally sent to CloudWatch in each account. The retention period for these logs should be three days.

Storage

The storage solution in the Log Archive account is implemented by using Amazon Simple Storage Service (Amazon S3). AWS Control Tower automatically sets up and manages the S3 buckets for AWS Control Tower according to AWS best practices.

The following table summarizes the storage configurations that you can configure in your landing zone. You should extend this table with additional storing solutions according to your landing zone requirements.

Account	S3 bucket name	Descripti on	Encryption	Lifecycle rules	Bucket policy	Created by AWS Control Tower?
Log Archive	`aws-contr`	This bucket is created	Default encryptio	The default	Default bucket	Yes

Account	S3 bucket name	Description	Encryption	Lifecycle rules	Bucket policy	Created by AWS Control Tower?
	`oltower-logs-*`	by AWS Control Tower and centralizes all AWS CloudTrail and AWS Config logs from all member accounts in your organization. Inside the bucket, files are kept in subdirectories that use the same account ID as the directory name.	n using SSE-S3 (AES-256)	retention period is 1 year. You can use AWS Control Tower customized log retention to extend log retention up to 15 years.	policy is applied.	

Account	S3 bucket name	Description	Encryption	Lifecycle rules	Bucket policy	Created by AWS Control Tower?
Log Archive	`aws-controltower-s3-access-logs-*`	This bucket is created by AWS Control Tower and collects the access logs of the first `aws-controltower-logs-*` S3 bucket.	Default encryption using SSE-S3 (AES-256)	The default retention period is 10 years. You can use AWS Control Tower customized log retention to extend log retention up to 15 years.	Default bucket policy is applied.	Yes
Shared Services	`aws-shared-services`	This S3 bucket is used to store the Amazon Machine Images (AMIs) for the landing zone.	Encryption using SSE-S3 (AES-256)	None.	Only accounts in the organization have access.	No

Encryption

Encryption is automatically enabled during landing zone setup for the S3 buckets that contain AWS Control Tower logs and access logs.

The S3 buckets for centralized logs should be encrypted at rest by using server-side encryption with Amazon S3 managed keys (SSE-S3). This option encrypts each object with a unique key by using 256-bit Advanced Encryption Standard (AES-256) encryption. As an additional safeguard, Amazon S3 encrypts the key itself with a management key that it regularly rotates.

You can also use server-side encryption with AWS Key Management Service (AWS KMS) keys. For more information, see the *Server-side encryption with AWS KMS keys (SSE-KMS)* section of Protecting data using server-side encryption in the Amazon S3 documentation. To configure AWS Control Tower to use a customer managed key (instead of the default AWS managed key), review the section Optionally configure AWS KMS keys in the AWS Control Tower documentation.

Auditing and alerting

The Audit account is tailored for auditors and security administrators. In this account, you can give auditors read-only access to all accounts in the organization, so they can conduct thorough reviews. Additionally, the Audit account can be the delegated administrator for several security services that monitor the accounts in the organization for threats and compliance.

Centralizing auditing and security services in a central AWS account offers numerous benefits, including:

- It isolates security functions from production workloads, to help collectively ensure robust and efficient security, compliance, and resource management across the organization's AWS environment.
- It simplifies visibility, security management, and incident response from one central place.
- It provides cost efficiency by eliminating redundancies.
- It enables automated remediations and alerts.

> **Note**
>
> When you set up alerts, you should also consider automating remediation actions by using AWS Config Rules, AWS Lambda functions, and AWS Systems Manager Automation documents.

The following table shows a recommended list of services to help manage and secure your landing zone. You should extend this table with additional monitoring solutions according to your landing zone requirements. For more guidance on security tooling you can include in the Audit account, see the AWS Security Reference Architecture.

Type	Description	Monitoring setup	Notification setup
Control compliance notifications	Provides notifications when there is drift in AWS Control Tower control compliance.	AWS Control Tower has an `aws-controltower-AggregateSecurityNotifications` SNS topic in the Audit account.	You should set up notifications after you create the AWS Control Tower landing zone to ensure that you can catch controls that are not compliant and in need of remediation. Note: You can automatically remediate non-compliant resources by using AWS Config Rules.
Threat detection (Amazon GuardDuty)	Monitors VPC Flow Logs, CloudTrail, and DNS logs to detect suspicious or unexpected behavior	We recommend that you set up and configure GuardDuty when you create the landing zone.	You should set up notifications after setting up GuardDuty to ensure that you receive alerts for

Type	Description	Monitoring setup	Notification setup
	in the accounts (for example, backdoor access, trojan programs, or unauthorized access). For more informati on, see the Amazon GuardDuty documentation.		potential threats to remediate. Note: You can integrate GuardDuty findings with AWS Security Hub.
Security and compliance monitorin g (AWS Security Hub)	Brings together security findings from multiple AWS services and third-party sources into a single centralized dashboard to help proactively identify and address security issues, vulnerabi lities, and compliance concerns. For more information, see the AWS Security Hub documentation.	We recommend that you set up and configure Security Hub when you create the landing zone.	You should set up notifications after setting up Security Hub to ensure that you receive alerts for potential vulnerabi lities to remediate. Note: You can automate remediati on in Security Hub.

Type	Description	Monitoring setup	Notification setup
Root user activity	Sends notifications when an account is accessed by the root user through the AWS Management Console.	We recommend that you set up an Amazon CloudWatch Events rule to monitor the `userIdentity` element in CloudTrail for root logins.	If there is root user account activity, CloudWatch Events writes to an SNS topic. For more information and an AWS CloudFormation script that you can use to set up this monitoring, see How do I create an EventBridge event rule to notify me that my AWS root user account was used? in the AWS Knowledge Center.

Type	Description	Monitoring setup	Notification setup
Billing alerts	Sends billing alerts if the cost and usage of AWS services exceeds your budget threshold.	We recommend that you set up a monthly customized budget that specifies a threshold that can be tracked by AWS Budgets.	AWS Budgets generates an alert by using Amazon Simple Notification Service (Amazon SNS) if the budget threshold is exceeded. You can use AWS CloudFormation stacks and an AWS CloudFormation template to set notifications at the organization or OU level. You can also choose to automatically apply this check to new accounts. For more information, see the AWS::Budgets::Budget resource in the AWS CloudFormation documentation.

Note

You can configure Amazon SNS to send out security alerts from the services listed in the table. The alerts can be sent to either one centralized email (if you have one single security team responsible), or to multiple emails (if different parts of your security organization are responsible for different services).

Managing the configuration of AWS resources

The AWS Config service enables you to assess, audit, and evaluate the configurations of your AWS resources. It provides a detailed view of how your resources are configured, shows how they relate to one another, and tracks how these configurations change over time. It's similar to a configuration management database that continuously monitors and records your AWS resource configurations, making it easier to audit resource compliance, analyze security postures, and troubleshoot configuration changes across your AWS environment. This service helps you maintain security and governance by tracking resource inventory, configuration history, and configuration change notifications to enable security and regulatory compliance.

Track resource configuration changes

AWS Control Tower enables AWS Config configuration recorders in all enrolled accounts to track resource configuration changes. For landing zone versions 3.0 and later, global resources (such as IAM users, groups, roles, and customer-managed policies) are recorded only in the home Region. For landing zone versions earlier than 3.0, these global resources are recorded in all enabled Regions. Each AWS Config recorder is set up with a delivery channel that sends all configuration changes to a centralized Amazon S3 bucket in the Log Archive account. This provides comprehensive tracking of resource configuration changes across the organization. For more information about how AWS Control Tower monitors resource changes with AWS Config, see Monitor resource changes with AWS Config in the AWS Control Tower documentation.

View configuration and compliance data

AWS Config aggregators provide a centralized way to view configuration and compliance data from multiple AWS accounts and Regions. They act as a central collector that consolidates AWS Config data across your organization and makes it easier to monitor resource configurations and compliance at scale. This capability is particularly valuable for enterprises that manage multiple AWS accounts, because it enables centralized auditing, governance, and compliance monitoring across their entire AWS footprint. AWS Control Tower creates two AWS Config aggregators to help manage and monitor your multi-account environment:

- **Organization-level aggregator** (`aws-controltower-ConfigAggregatorForOrganizations`) is created in the management account of your AWS organization. Its primary purpose is to aggregate AWS Config data from all accounts in your

organization, even if those accounts aren't enrolled in AWS Control Tower. AWS Config isn't enabled in the managed account by default, so you can't see this aggregator in the AWS Config console. To view the aggregator in the management account, use the AWS CLI command:

```
aws configservice describe-configuration-aggregators
```

- **Security aggregator** (`aws-controltower-GuardRailsComplianceAggregator`) is created in the Audit account of your AWS Control Tower environment. Its primary purpose is to monitor compliance with AWS Control Tower guardrails. It aggregates the relevant AWS Config data from all accounts that are enrolled in AWS Control Tower.

FAQ

This section provides answers to commonly raised questions about designing an AWS Control Tower landing zone.

Q. How can I use this guide in my design document?

A. You can replicate different sections in this guide and use it as a template for your design document. You can also adapt the tables and diagrams in this guide to your specific requirements and configurations.

Q. Where can I find more information about AWS Control Tower?

A. For sample features, FAQs, and pricing, see the AWS Control Tower web pages. For the user guide, best practices, and tutorials, see the AWS Control Tower documentation.

Q. How can I get started with AWS Control Tower?

A. Use the AWS Control Tower Workshop to set up your AWS Control Tower landing zone.

Q. Where can I find architectures for AWS Control Tower?

A.Visit the AWS Architecture Blog, which has an AWS Control Tower category that provides different architectural solutions and customizations.

Q. How can I further customize my AWS Control Tower landing zone?

A. You can use the Customizations for AWS Control Tower or Landing Zone Accelerator solution to customize your AWS Control Tower landing zone. To customize the accounts provisioned, you can use the Account Factory Customization (AFC) or the Account Factory for Terraform (AFT) solution. For additional solutions that you can integrate with AWS Control Tower, see AWS Control Tower in AWS Marketplace.

Q. Where can I find up-to-date information about all AWS Control Tower controls?

A. For complete information about current controls, see Controls reference guide in the AWS Control Tower documentation.

Q. How can I automate operations in AWS Control Tower?

A. AWS Control Tower provides an API that exposes some of the most commonly used operations such as enabling or disabling controls, or provisioning or closing accounts. For more information, see Automate tasks in AWS Control Tower in the AWS Control Tower documentation.

Q. Where can I find more best practices for architecting multi-account environments on AWS?

A. See the white paper Organizing Your AWS Environment Using Multiple Accounts for additional guidance.

Q. Where can I find more best practices for designing a secure landing zone?

A. See the AWS Security Reference Architecture for best practices for building a secure landing zone.

Resources

- AWS Control Tower documentation
- Setting up (AWS Control Tower documentation)
- Controls reference (AWS Control Tower documentation)
- How does AWS Control Tower establish your multi-account environment? (AWS whitepaper)
- AWS multi-account strategy for your AWS Control Tower landing zone (AWS Control Tower documentation)
- Best practices for AWS Control Tower administrators (AWS Control Tower documentation)
- AWS Control Tower Account Factory for Terraform (AFT) (AWS Control Tower documentation)
- AWS Cloud Adoption Framework
- AWS Well-Architected Framework
- Hybrid Networking Lens (AWS Well-Architected Framework)
- Migrate an AWS member account from AWS Organizations to AWS Control Tower (AWS Prescriptive Guidance)
- Organizing Your AWS Environment Using Multiple Accounts (AWS whitepaper)
- AWS Security Reference Architecture (AWS Prescriptive Guidance)
- Setting up a secure and scalable multi-account AWS environment (AWS Prescriptive Guidance)

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
New section	Added information about managing the configuration of AWS resources; updated AWS Direct Connect details in the Networking integration section.	December 6, 2024
Major update	Significant changes and additions to all sections of the guide.	March 25, 2024
Update	Updated with the latest details on AWS Control Tower controls.	November 17, 2022
—	Initial publication	September 10, 2021

AWS Prescriptive Guidance glossary

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

Numbers

7 Rs

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

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

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

A

ABAC

See attribute-based access control.

abstracted services

See managed services.

ACID

See atomicity, consistency, isolation, durability.

active-active migration

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

active-passive migration

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

aggregate function

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

AI

See artificial intelligence.

AIOps

See artificial intelligence operations.

anonymization

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

anti-pattern

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

application control

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

application portfolio

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

artificial intelligence (AI)

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

artificial intelligence operations (AIOps)

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

asymmetric encryption

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

atomicity, consistency, isolation, durability (ACID)

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

attribute-based access control (ABAC)

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

authoritative data source

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

Availability Zone

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

AWS Cloud Adoption Framework (AWS CAF)

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

AWS Workload Qualification Framework (AWS WQF)

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

B

bad bot

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

BCP

See business continuity planning.

behavior graph

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

big-endian system

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

binary classification

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

bloom filter

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

blue/green deployment

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

bot

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

botnet

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

branch

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

break-glass access

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

brownfield strategy

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

buffer cache

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

business capability

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

business continuity planning (BCP)

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

C

CAF

See AWS Cloud Adoption Framework.

canary deployment

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

CCoE

See Cloud Center of Excellence.

CDC

See change data capture.

change data capture (CDC)

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

chaos engineering

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

CI/CD

See continuous integration and continuous delivery.

classification

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

client-side encryption

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

Cloud Center of Excellence (CCoE)

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

cloud computing

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

cloud operating model

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

cloud stages of adoption

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

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

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

CMDB

See configuration management database.

code repository

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

cold cache

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

cold data

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

computer vision (CV)

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

configuration drift

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

configuration management database (CMDB)

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

conformance pack

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

continuous integration and continuous delivery (CI/CD)

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

CV

See computer vision.

D

data at rest

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

data classification

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

data drift

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

data in transit

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

data mesh

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

data minimization

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

data perimeter

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

data preprocessing

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

data provenance

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

data subject

An individual whose data is being collected and processed.

data warehouse

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

database definition language (DDL)

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

database manipulation language (DML)

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

DDL

See database definition language.

deep ensemble

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

deep learning

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

defense-in-depth

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

delegated administrator

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

deployment

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

development environment

See environment.

detective control

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

development value stream mapping (DVSM)

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

digital twin

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

dimension table

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

disaster

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

disaster recovery (DR)

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

DML

See database manipulation language.

domain-driven design

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

DR

See disaster recovery.

drift detection

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

DVSM

See development value stream mapping.

E

EDA

See exploratory data analysis.

EDI

See electronic data interchange.

edge computing

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

electronic data interchange (EDI)

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

encryption

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

encryption key

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

endianness

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

endpoint

See service endpoint.

endpoint service

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

information, see Create an endpoint service in the Amazon Virtual Private Cloud (Amazon VPC) documentation.

enterprise resource planning (ERP)

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

envelope encryption

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

environment

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

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

epic

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

ERP

See enterprise resource planning.

exploratory data analysis (EDA)

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

F

fact table

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

fail fast

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

fault isolation boundary

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

feature branch

See branch.

features

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

feature importance

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

feature transformation

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

few-shot prompting

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

FGAC

See fine-grained access control.

fine-grained access control (FGAC)

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

flash-cut migration

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

FM

See foundation model.

foundation model (FM)

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

G

generative AI

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

geo blocking

See geographic restrictions.

geographic restrictions (geo blocking)

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

Gitflow workflow

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

golden image

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

greenfield strategy

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

guardrail

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

Detective guardrails detect policy violations and compliance issues, and generate alerts for remediation. They are implemented by using AWS Config, AWS Security Hub, Amazon GuardDuty, AWS Trusted Advisor, Amazon Inspector, and custom AWS Lambda checks.

H

HA

See high availability.

heterogeneous database migration

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

high availability (HA)

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

historian modernization

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

holdout data

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

homogeneous database migration

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

hot data

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

hotfix

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

hypercare period

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

I

IaC

See infrastructure as code.

identity-based policy

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

idle application

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

IIoT

See industrial Internet of Things.

immutable infrastructure

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

inbound (ingress) VPC

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

incremental migration

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

Industry 4.0

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

infrastructure

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

infrastructure as code (IaC)

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

industrial Internet of Things (IIoT)

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

inspection VPC

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

Internet of Things (IoT)

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

interpretability

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

IoT

See Internet of Things.

IT information library (ITIL)

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

IT service management (ITSM)

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

ITIL

See IT information library.

ITSM

See IT service management.

L

label-based access control (LBAC)

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

landing zone

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

large language model (LLM)

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

large migration

A migration of 300 or more servers.

LBAC

See label-based access control.

least privilege

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

lift and shift

See 7 Rs.

little-endian system

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

LLM

See large language model.

lower environments

See environment.

M

machine learning (ML)

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

main branch

See branch.

malware

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

managed services

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

manufacturing execution system (MES)

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

MAP

See Migration Acceleration Program.

mechanism

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

member account

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

MES

See manufacturing execution system.

Message Queuing Telemetry Transport (MQTT)

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

microservice

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

microservices architecture

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

Migration Acceleration Program (MAP)

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

migration at scale

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

migration factory

Cross-functional teams that streamline the migration of workloads through automated, agile approaches. Migration factory teams typically include operations, business analysts and owners,

migration engineers, developers, and DevOps professionals working in sprints. Between 20 and 50 percent of an enterprise application portfolio consists of repeated patterns that can be optimized by a factory approach. For more information, see the discussion of migration factories and the Cloud Migration Factory guide in this content set.

migration metadata

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

migration pattern

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

Migration Portfolio Assessment (MPA)

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

Migration Readiness Assessment (MRA)

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

migration strategy

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

ML

See machine learning.

modernization

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

modernization readiness assessment

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

monolithic applications (monoliths)

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

MPA

See Migration Portfolio Assessment.

MQTT

See Message Queuing Telemetry Transport.

multiclass classification

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

mutable infrastructure

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

O

OAC

See origin access control.

OAI

See origin access identity.

OCM

See organizational change management.

offline migration

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

OI

See operations integration.

OLA

See operational-level agreement.

online migration

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

OPC-UA

See Open Process Communications - Unified Architecture.

Open Process Communications - Unified Architecture (OPC-UA)

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

operational-level agreement (OLA)

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

operational readiness review (ORR)

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

operational technology (OT)

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

operations integration (OI)

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

organization trail

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

organizational change management (OCM)

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

origin access control (OAC)

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

origin access identity (OAI)

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

ORR

See operational readiness review.

OT

See operational technology.

outbound (egress) VPC

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

P

permissions boundary

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

personally identifiable information (PII)

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

PII

See personally identifiable information.

playbook

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

PLC

See programmable logic controller.

PLM

See product lifecycle management.

policy

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

polyglot persistence

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

portfolio assessment

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

predicate

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

predicate pushdown

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

preventative control

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

principal

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

privacy by design

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

private hosted zones

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

proactive control

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

product lifecycle management (PLM)

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

production environment

See environment.

programmable logic controller (PLC)

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

prompt chaining

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

pseudonymization

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

publish/subscribe (pub/sub)

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

Q

query plan

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

query plan regression

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

R

RACI matrix

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

RAG

See Retrieval Augmented Generation.

ransomware

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

RASCI matrix

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

RCAC

See row and column access control.

read replica

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

re-architect

See 7 Rs.

recovery point objective (RPO)

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

recovery time objective (RTO)

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

refactor

See 7 Rs.

Region

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

regression

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

rehost

See 7 Rs.

release

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

relocate

See 7 Rs.

replatform

See 7 Rs.

repurchase

See 7 Rs.

resiliency

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

resource-based policy

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

responsible, accountable, consulted, informed (RACI) matrix

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

responsive control

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

retain

See 7 Rs.

retire

See 7 Rs.

Retrieval Augmented Generation (RAG)

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

rotation

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

row and column access control (RCAC)

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

RPO

See recovery point objective.

RTO

See recovery time objective.

runbook

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

S

SAML 2.0

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

SCADA

See supervisory control and data acquisition.

SCP

See service control policy.

secret

In AWS Secrets Manager, confidential or restricted information, such as a password or user credentials, that you store in encrypted form. It consists of the secret value and its metadata.

The secret value can be binary, a single string, or multiple strings. For more information, see What's in a Secrets Manager secret? in the Secrets Manager documentation.

security by design

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

security control

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

security hardening

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

security information and event management (SIEM) system

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

security response automation

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

server-side encryption

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

service control policy (SCP)

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

service endpoint

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

service-level agreement (SLA)

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

service-level indicator (SLI)

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

service-level objective (SLO)

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

shared responsibility model

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

SIEM

See security information and event management system.

single point of failure (SPOF)

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

SLA

See service-level agreement.

SLI

See service-level indicator.

SLO

See service-level objective.

split-and-seed model

A pattern for scaling and accelerating modernization projects. As new features and product releases are defined, the core team splits up to create new product teams. This helps scale your

organization's capabilities and services, improves developer productivity, and supports rapid innovation. For more information, see Phased approach to modernizing applications in the AWS Cloud.

SPOF

See single point of failure.

star schema

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

strangler fig pattern

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

subnet

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

supervisory control and data acquisition (SCADA)

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

symmetric encryption

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

synthetic testing

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

system prompt

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

T

tags

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

target variable

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

task list

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

test environment

See environment.

training

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

transit gateway

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

trunk-based workflow

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

trusted access

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

tuning

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

two-pizza team

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

U

uncertainty

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

undifferentiated tasks

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

upper environments

See environment.

V

vacuuming

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

version control

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

VPC peering

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

vulnerability

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

W

warm cache

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

warm data

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

window function

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

workload

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

workstream

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

WORM

See write once, read many.

WQF

See AWS Workload Qualification Framework.

write once, read many (WORM)

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

Z

zero-day exploit

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

zero-day vulnerability

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

zero-shot prompting

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

zombie application

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