
AWS Prescriptive Guidance

AWS Startup Security Baseline (AWS SSB)



AWS Prescriptive Guidance: AWS Startup Security Baseline (AWS SSB)

Copyright © 2023 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

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

Table of Contents

Introduction	1
Intended audience	1
Foundational framework and security responsibilities	1
Securing your account	3
ACCT.01 – Set account-level contacts	3
ACCT.02 – Restrict use of the root user	4
ACCT.03 – Configure console access	4
ACCT.04 – Assign permissions	5
ACCT.05 – Require MFA	5
ACCT.06 – Enforce a password policy	6
ACCT.07 – Log events	7
ACCT.08 – Prevent public access to private S3 buckets	7
ACCT.09 – Delete unused resources	8
ACCT.10 – Monitor costs	8
ACCT.11 – Enable GuardDuty	9
ACCT.12 – Monitor high-risk issues	9
Securing your workloads	10
WKLD.01 – Use IAM roles for permissions	10
WKLD.02 – Use resource-based policies	11
WKLD.03 – Use ephemeral secrets or a secrets-management service	11
WKLD.04 – Protect application secrets	12
WKLD.05 – Detect and remediate exposed secrets	13
WKLD.06 – Use Systems Manager instead of SSH or RDP	13
WKLD.07 – Log data events for select S3 buckets	14
WKLD.08 – Encrypt Amazon EBS volumes	14
WKLD.09 – Encrypt Amazon RDS databases	14
WKLD.10 – Deploy private resources in private subnets	15
WKLD.11 – Use security groups to restrict access	15
WKLD.12 – Use VPC endpoints to access services	16
WKLD.13 – Require HTTPS for all public web endpoints	17
WKLD.14 – Use edge-protection services for public endpoints	18
WKLD.15 – Use templates to deploy security controls	18
Contributors	19
Document history	20
Glossary	21
Security terms	21

AWS Startup Security Baseline (AWS SSB)

Jay Michael, Amazon Web Services (AWS)

May 2023 ([document history \(p. 20\)](#))

The AWS Startup Security Baseline (SSB) is a set of controls that create a minimum foundation for businesses to build securely on AWS without decreasing their agility. These controls form the basis of your security posture and are focused on securing credentials, enabling logging and visibility, managing contact information, and implementing basic data boundaries.

The controls in this guide are designed with early startups in mind, mitigating the most common security risks without requiring significant effort. Many startups begin their journey in the AWS Cloud with a single AWS account. As organizations grow, they migrate to multi-account architectures. The guidance in this guide is designed for single-account architectures, but it helps you set up security controls that are easily migrated or modified as you transition to a multi-account architecture.

The controls in the AWS SSB are separated into two categories: account and workload. Account controls help keep your AWS account secure. It includes recommendations for setting up user access, policies, and permissions, and it includes recommendations for how to monitor your account for unauthorized or potentially malicious activity. Workload controls help secure your resources and code in the cloud, such as applications, backend processes, and data. It includes recommendations such as encryption and reducing the scope of access.

Note

Some of the controls recommended in this guide replace the defaults configured during initial setup, while most configure new settings and policies. This document should in no way be considered comprehensive of all available controls.

Intended audience

This guide is best suited for startups that are in the very beginning stages of development, with minimal staff and operations.

Startups or other businesses that are in later stages of operation and growth can still derive significant value from reviewing these controls against their current practices. If you identify any gaps, you can implement the individual controls in this guide and then evaluate them for appropriateness as a long-term solution.

Note

The recommended controls in this guide are foundational in nature. Startups or other companies operating at a later stage of scale or sophistication should add additional controls as applicable.

Foundational framework and security responsibilities

[AWS Well-Architected](#) helps cloud architects build a secure, high-performing, resilient, and efficient infrastructure for their applications and workloads. The AWS Startup Security Baseline aligns to

the [security pillar](#) of the AWS Well-Architected Framework. The *security pillar* describes how to take advantage of cloud technologies to protect data, systems, and assets in a way that can improve your security posture. This helps you meet your business and regulatory requirements by following current AWS recommendations.

You can assess your adherence to Well-Architected best practices by using the [AWS Well-Architected Tool](#) in your AWS account.

Security and compliance are a shared responsibility between AWS and the customer. The [shared responsibility model](#) is often described by saying that AWS is responsible for the security of the cloud (that is, for protecting the infrastructure that runs all the services offered in the AWS Cloud), and you are responsible for the security in the cloud (as determined by the AWS Cloud services that you select). In the shared responsibility model, implementing the security controls in this document is part of your responsibility as a customer.

Securing your account

Controls and recommendations in this section help keep your AWS account secure. It emphasizes using AWS Identity and Access Management (IAM) users, user groups, and roles (also known as *principals*) for both human and machine access, restricting the use of the root user, and requiring multi-factor authentication. In this section, you confirm that AWS has the contact information necessary to reach you regarding your account activity and status. You also set up monitoring services, such as AWS Trusted Advisor, Amazon GuardDuty, and AWS Budgets, so that you are notified of activity in your account and can respond quickly if the activity is unauthorized or unexpected.

This section contains the following topics:

- [ACCT.01 – Set account-level contacts to valid email distribution lists \(p. 3\)](#)
- [ACCT.02 – Restrict use of the root user \(p. 4\)](#)
- [ACCT.03 – Configure console access for each user \(p. 4\)](#)
- [ACCT.04 – Assign permissions \(p. 5\)](#)
- [ACCT.05 – Require multi-factor authentication \(MFA\) to log in \(p. 5\)](#)
- [ACCT.06 – Enforce a password policy \(p. 6\)](#)
- [ACCT.07 – Deliver CloudTrail logs to a protected S3 bucket \(p. 7\)](#)
- [ACCT.08 – Prevent public access to private S3 buckets \(p. 7\)](#)
- [ACCT.09 – Delete unused VPCs, subnets, and security groups \(p. 8\)](#)
- [ACCT.10 – Configure AWS Budgets to monitor your spending \(p. 8\)](#)
- [ACCT.11 – Enable and respond to GuardDuty notifications \(p. 9\)](#)
- [ACCT.12 – Monitor for and resolve high-risk issues by using Trusted Advisor \(p. 9\)](#)

ACCT.01 – Set account-level contacts to valid email distribution lists

When setting up primary and alternate contacts for your AWS account, use an email distribution list instead of an individual's email address. Using an email distribution list makes sure that ownership and reachability are preserved as individuals in your organization come and go. Set alternate contacts for billing, operations, and security notifications, and use appropriate email distribution lists accordingly. AWS uses these email addresses to contact you, so it is important you retain access to them.

To edit your account name, root user password, or root user email address

1. Sign in to the **Account Settings** page in the Billing and Cost Management console at <https://console.aws.amazon.com/billing/home?#/account>.
2. On the **Account Settings** page, next to **Account Settings**, choose **Edit**.
3. Next to the field you want to update, choose **Edit**.
4. After you have entered your changes, choose **Save changes**.
5. After you have made all of your changes, choose **Done**.

To edit your contact information

1. On the [Account Settings](#) page, under **Contact Information**, choose **Edit**.

2. For the fields you want to change, enter your updated information, and then choose **Update**.

To add, update, or remove alternate contacts

1. On the [Account Settings](#) page, under **Alternate Contacts**, choose **Edit**.
2. For the fields you want to change, enter your updated information, and then choose **Update**.

ACCT.02 – Restrict use of the root user

The root user is created when you sign up for an AWS account, and this user has full ownership privileges and permissions over the account that cannot be changed. Only use the root user for the specific tasks that require it. For more information, see [Tasks that require root user credentials](#) (AWS Account Management). Perform all other actions in your account by using other types of IAM identities, such as federated users with IAM roles. For more information, see [AWS security credentials](#) (IAM documentation).

To restrict use of the root user

1. Require multi-factor authentication (MFA) for the root user as described in [ACCT.05 – Require multi-factor authentication \(MFA\) to log in \(p. 5\)](#).
2. Create an administrative user so that you don't use the root user for everyday tasks. For more information about configuring user access, see [ACCT.03 – Configure console access for each user \(p. 4\)](#).

ACCT.03 – Configure console access for each user

As a best practice, AWS recommends using temporary credentials to grant access to AWS accounts and resources. *Temporary credentials* have a limited lifetime, so you do not have to rotate them or explicitly revoke them when they're no longer needed. For more information, see [Temporary security credentials](#) (IAM documentation).

For human users, AWS recommends using federated identities from a centralized identity provider (IdP), such as AWS IAM Identity Center, Okta, Active Directory, or Ping Identity. Federating users allows you to define identities in a single, central location, and users can securely authenticate to multiple applications and websites, including AWS, by using just one set of credentials. For more information, see [Identity federation in AWS](#) and [IAM Identity Center](#) (AWS website).

Note

Identity federation can complicate the transition from a single-account architecture to a multi-account architecture. It is common for startups to delay implementing identity federation until they have established a multi-account architecture managed in AWS Organizations.

To set up identity federation

1. If you are using IAM Identity Center, see [Getting started](#) (IAM Identity Center documentation).
If you are using an external or third-party IdP, see [Creating IAM identity providers](#) (IAM documentation).
2. Make sure that your IdP enforces multi-factor authentication (MFA).
3. Apply permissions according to [ACCT.04 – Assign permissions \(p. 5\)](#).

For startups that are not prepared to configure identity federation, you can create users directly in IAM. This is not a recommended security best practice because these are long-term credentials that never

expire. However, this is a common practice for startups in early operation to prevent difficulty with transitioning to a multi-account architecture when they're operationally ready.

As a baseline, you can create an IAM user for each person who needs to access the AWS Management Console. If you configure IAM users, do not share credentials across users, and regularly rotate the long-term credentials.

To create an IAM user

1. [Create IAM users](#) (IAM documentation).
2. Apply permissions according to [ACCT.04 – Assign permissions \(p. 5\)](#).

ACCT.04 – Assign permissions

Configure user permissions in the account by assigning *policies* to their IAM identity (user group or role). You can customize the permissions, or you can attach [AWS managed policies](#), which are standalone policies designed by AWS to provide permissions for many common use cases. If you customize permissions, follow the security best practice of [granting least privilege](#). *Least privilege* is the practice of granting the minimum set of permissions that each user needs to perform their tasks.

If you are using federated identities, users access the account by assuming an IAM role through the external identity provider. The IAM role defines what users authenticated by your organization's IdP are allowed to do in AWS. You apply custom or AWS managed policies to this role to configure permissions.

To assign permissions for federated identities

- If you are using IAM Identity Center, see [Use IAM policies in permission sets](#) (IAM Identity Center documentation).

If you are using an external or third-party IdP, see [Adding IAM identity permissions](#) (IAM documentation).

If you are using IAM users, you can use user groups or roles to manage permissions for multiple IAM users. We recommend user groups for startups because they are easier to manage and less prone to misconfiguration that could pose security risks for your account. Assign users to user groups based on their job functions. Examples of user groups include application, data, networking, and Development Operations (DevOps) engineers. You can also divide the user types into smaller user groups based on decision-making authority, such as for senior or non-senior engineers.

To assign permissions for IAM users

1. [Create IAM user groups](#) (IAM documentation).
2. [Attach an AWS managed policy to an IAM user group](#) (IAM documentation).

ACCT.05 – Require multi-factor authentication (MFA) to log in

With MFA, users have a device that generates a response to an authentication challenge. Each user's credentials and device-generated response are required to complete the sign-in process. As a security best practice, enable MFA for AWS account access, especially for long-term credentials such as the account root user and IAM users.

To set up MFA for the root user

1. Sign in to the AWS Management Console at <https://console.aws.amazon.com/>.
2. On the right side of the navigation bar, choose your account name, and then choose **My Security Credentials**.
3. If necessary, choose **Continue to Security Credentials**.
4. Expand the **Multi-Factor Authentication (MFA)** section.
5. Choose **Activate MFA**.
6. Follow the wizard instructions to configure your MFA devices accordingly. For more information, see [Enabling MFA devices for users in AWS](#) (IAM documentation).

To set up MFA in IAM Identity Center

- [Enable MFA](#) (IAM Identity Center documentation)

To set up MFA for your own IAM user

1. Using your sign-in credentials, sign in to the IAM console at <https://console.aws.amazon.com/iam>.
2. In the navigation bar on the upper right, choose your user name, and then choose **My Security Credentials**.
3. On the **AWS IAM credentials** tab, in the **Multi-factor authentication** section, choose **Manage MFA device**.

To set up MFA for other IAM users

1. Sign in to the AWS Management Console and open the IAM console at <https://console.aws.amazon.com/iam>.
2. In the navigation pane, choose **Users**.
3. Choose the name of the user for whom you want to enable MFA, and then choose the **Security credentials** tab.
4. Next to **Assigned MFA device**, choose **Manage**.
5. Follow the wizard instructions to configure your MFA devices accordingly. For more information, see [Enabling MFA devices for users in AWS](#) (IAM documentation).

ACCT.06 – Enforce a password policy

Users log in to the AWS Management Console by providing sign-in credentials, and MFA is recommended. Require that passwords adhere to a strong password policy to help prevent discovery through brute force or social engineering.

For more information about the latest recommendations for strong passwords, see [Password Policy Guide](#) on the Center for Internet Security (CIS) website.

For IAM users, you can configure password requirements in a custom IAM password policy. For more information, see [Setting an account password policy](#) (IAM documentation).

To create a custom password policy

1. Sign in to the AWS Management Console and open the IAM console at <https://console.aws.amazon.com/iam>.

2. In the navigation pane, choose **Account settings**.
3. In the **Password policy** section, choose **Change password policy**.
4. Select the options that you want to apply to your password policy, and then choose **Save changes**.

ACCT.07 – Deliver CloudTrail logs to a protected S3 bucket

Actions taken by users, roles, and services in your AWS account are recorded as events in AWS CloudTrail. CloudTrail is enabled by default, and in the CloudTrail console, you can access 90 days of event history information. To view, search, download, archive, analyze, and respond to account activity across your AWS infrastructure, see [Viewing events with CloudTrail Event history](#) (CloudTrail documentation).

To retain CloudTrail history beyond 90 days with additional data, you create a new trail that delivers log files to an Amazon Simple Storage Service (Amazon S3) bucket for all event types. When you create a trail in the CloudTrail console, you create a multi-region trail.

To create a trail that delivers logs for all AWS Regions to an S3 bucket

1. [Create a trail](#) (CloudTrail documentation). On the **Choose log events** page, do the following:
 - a. For **API activity**, choose both **Read** and **Write**.
 - b. For preproduction environments, choose **Exclude AWS KMS events**. This excludes all AWS Key Management Service (AWS KMS) events from your trail. AWS KMS **read** actions such as `Encrypt`, `Decrypt`, and `GenerateDataKey` can generate a large volume of events.

For production environments, choose to log **Write** management events, and clear the check box for **Exclude AWS KMS events**. This excludes high-volume AWS KMS read events but still logs relevant write events, such as `Disable`, `Delete`, and `ScheduleKey`. These are the minimum recommended AWS KMS logging settings for a production environment.

2. The new trail appears on the **Trails** page. In about 15 minutes, CloudTrail publishes log files that show the AWS application programming interface (API) calls made in your account. You can see the log files in the S3 bucket that you specified.

To help secure the S3 buckets where you store CloudTrail log files

1. Review the [Amazon S3 bucket policy](#) (CloudTrail documentation) for any and all buckets where you store log files and adjust it as needed to remove any unnecessary access.
2. As a security best practice, be sure to manually add an `aws:SourceArn` condition key to the bucket policy. For more information, see [Create or update an Amazon S3 bucket to use to store the log files for an organization trail](#) (CloudTrail documentation).
3. [Enable MFA Delete](#) (Amazon S3 documentation).

ACCT.08 – Prevent public access to private S3 buckets

By default, only the root user of the AWS account and the IAM principal, if used, have permissions to read and write to Amazon S3 buckets created by that principal. Additional IAM principals are granted access by using identity-based policies, and access conditions can be enforced by using a bucket policy. You can create bucket policies that grant the general public access to the bucket, a *public* bucket.

Buckets created on or after April 28, 2023 have the **Block Public Access** setting enabled by default. For buckets created before this date, users might misconfigure the bucket policy and unintentionally grant access to the public. You can prevent this misconfiguration by enabling the **Block Public Access** setting for each bucket. If you have no current or future use cases for a public S3 bucket, enable this setting at the AWS account level. This setting prevents policies that allow public access.

To prevent public access to S3 buckets

- [Configure block public access settings for your S3 buckets](#) (Amazon S3 documentation).

AWS Trusted Advisor generates a yellow finding for S3 buckets that allow list or read access to the public and generates a red finding for buckets that allow public uploads or deletes. As a baseline, follow the control [ACCT.12 – Monitor for and resolve high-risk issues by using Trusted Advisor \(p. 9\)](#) to identify and correct misconfigured buckets. Publicly accessible S3 buckets are also indicated in the Amazon S3 console.

ACCT.09 – Delete unused VPCs, subnets, and security groups

To reduce the opportunity for security issues, delete or turn off any resources that are not being used. In a new AWS account, by default a virtual private cloud (VPC) is created automatically in every AWS Region, which enables you to assign public IP addresses in public subnets. However, if these VPCs are not needed, this introduces risk of unintended exposure of resources.

If they are not in use, delete the default VPCs in all Regions, not just those in the Regions where you might deploy workloads. Deleting a VPC also deletes its components, such as subnets and security groups.

Note

You can view all Regions and VPCs on the Amazon EC2 Global View console at <https://console.aws.amazon.com/ec2globalview/home>. For more information, see [List and filter resources across Regions using Amazon EC2 Global View](#) (Amazon EC2 documentation).

To delete unused default VPCs

1. [Delete your VPC](#) (Amazon VPC documentation).
2. Repeat as needed for VPCs in other Regions.

ACCT.10 – Configure AWS Budgets to monitor your spending

AWS Budgets enable monitoring of monthly costs and usage with notifications when costs are forecasted to exceed target thresholds. Forecasted cost notifications can provide an indication of unexpected activity, providing extra defense in addition to other monitoring systems, such as AWS Trusted Advisor and Amazon GuardDuty. Monitoring and understanding your AWS costs is also part of good operational hygiene.

To set up a budget in AWS Budgets

- [Create a cost budget](#) (AWS Budgets documentation).

ACCT.11 – Enable and respond to GuardDuty notifications

Amazon GuardDuty is a threat-detection service that continuously monitors for malicious or unauthorized behavior to help protect your AWS accounts, workloads, and data. When it detects unexpected and potentially malicious activity, GuardDuty delivers detailed security findings for visibility and remediation. GuardDuty can detect threats such as cryptocurrency mining activity, access from Tor clients and relays, unexpected behavior, and compromised IAM credentials. Enable GuardDuty and respond to findings to stop potentially malicious or unauthorized behavior in your AWS environment. For more information about findings in GuardDuty, see [Finding types](#) (GuardDuty documentation).

You can use Amazon CloudWatch Events to set up automated notifications when GuardDuty creates a finding or the finding changes. First, you set up an Amazon Simple Notification Service (Amazon SNS) topic and add endpoints, or email addresses, to the topic. Then, you set up a CloudWatch event for GuardDuty findings, and the event rule notifies the endpoints in the Amazon SNS topic.

To enable GuardDuty and GuardDuty notifications

1. [Enable Amazon GuardDuty](#) (GuardDuty documentation).
2. [Create a CloudWatch Events rule to notify you of GuardDuty findings](#) (GuardDuty documentation).

ACCT.12 – Monitor for and resolve high-risk issues by using Trusted Advisor

AWS Trusted Advisor passively scans your AWS infrastructure for high-risk or high-impact issues related to security, performance, cost, and reliability. It provides detailed information about affected resources and remediation recommendations. For a complete list of checks and descriptions, see [AWS Trusted Advisor check reference](#) (Trusted Advisor documentation).

Review Trusted Advisor findings on a recurring basis, and remediate issues as necessary. If you have the AWS Business Support or Enterprise Support plans, you can subscribe to a weekly findings email. For more information, see [Set up notification preferences](#) (AWS Support documentation).

To view issues in Trusted Advisor

- Review each check category according to the instructions in [View check categories](#) (AWS Support documentation). At a minimum, we recommend reviewing the **action recommended** issues, which are red.

Securing your workloads

Controls and recommendations in this section help you secure your workloads running in AWS, while you are building them. They emphasize secure practices for managing application secrets and scope of access, minimizing access routes to private resources, and using encryption to protect data in transit and at rest.

This section contains the following topics:

- [WKLD.01 – Use IAM roles for compute environment permissions \(p. 10\)](#)
- [WKLD.02 – Restrict credential usage scope with resource-based policies permissions \(p. 11\)](#)
- [WKLD.03 – Use ephemeral secrets or a secrets-management service \(p. 11\)](#)
- [WKLD.04 – Prevent application secrets from being exposed \(p. 12\)](#)
- [WKLD.05 – Detect and remediate exposed secrets \(p. 13\)](#)
- [WKLD.06 – Use Systems Manager instead of SSH or RDP \(p. 13\)](#)
- [WKLD.07 – Log data events for S3 buckets with sensitive data \(p. 14\)](#)
- [WKLD.08 – Encrypt Amazon EBS volumes \(p. 14\)](#)
- [WKLD.09 – Encrypt Amazon RDS databases \(p. 14\)](#)
- [WKLD.10 – Deploy private resources into private subnets \(p. 15\)](#)
- [WKLD.11 – Restrict network access by using security groups \(p. 15\)](#)
- [WKLD.12 – Use VPC endpoints to access supported services \(p. 16\)](#)
- [WKLD.13 – Require HTTPS for all public web endpoints \(p. 17\)](#)
- [WKLD.14 – Use edge-protection services for public endpoints \(p. 18\)](#)
- [WKLD.15 – Define security controls in templates and deploy them by using CI/CD practices \(p. 18\)](#)

WKLD.01 – Use IAM roles for compute environment permissions

In AWS Identity and Access Management (IAM), a *role* represents a set of permissions that can be assumed by a person or service for a configurable period of time. Using roles eliminates the need to store or manage long-term credentials, significantly reducing the chance of unintended use. Assign an IAM role directly to Amazon Elastic Compute Cloud (Amazon EC2) instances, AWS Fargate tasks and services, AWS Lambda functions, and other AWS compute services whenever supported. Applications that use an AWS SDK and run in these compute environments automatically use the IAM role credentials for authentication.

The approach and instructions for using IAM roles for each service can be found in the [AWS Documentation](#) for the service. For example, see the following:

- [IAM roles for Amazon EC2](#) (Amazon EC2 documentation)
- [IAM roles for tasks](#) (Amazon Elastic Container Service documentation)
- [Lambda execution role](#) (Lambda documentation)

WKLD.02 – Restrict credential usage scope with resource-based policies permissions

Policies are objects that can define permissions or specify access conditions. There are two primary types of policies:

- *Identity-based policies* are attached to principals and define what the principal's permissions in the AWS environment.
- *Resource-based policies* are attached to a resource, such as an Amazon Simple Storage Service (Amazon S3) bucket, or virtual private cloud (VPC) endpoint. These policies specify which principals are allowed access, supported actions, and any other conditions that must be met.

For a principal to be allowed access to perform an action against a resource, it must have permission granted in its identity-based policy and meet the conditions of the resource-based policy. For more information, see [Identity-based policies and resource-based policies](#) (IAM documentation).

Recommended conditions for resource-based policies include:

- Restrict access to only principals in a specified organization (defined in AWS Organizations) by using the `aws:PrincipalOrgID` condition.
- Restrict access to traffic that originates from a specific VPC or VPC endpoint by using the `aws:SourceVpc` or `aws:SourceVpce` condition, respectively.
- Allow or deny traffic based on the source IP address by using an `aws:SourceIp` condition.

The following is an example of a resource-based policy that uses the `aws:PrincipalOrgID` condition to allow only principals in the `<o-xxxxxxxxxxx>` organization to access the `<bucket-name>` S3 bucket:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowFromOrganization",
      "Effect": "Allow",
      "Principal": "*",
      "Action": "s3:*",
      "Resource": "arn:aws:s3:::<bucket-name>/*",
      "Condition": {
        "StringEquals": {"aws:PrincipalOrgID": "<o-xxxxxxxxxxx>"}
      }
    }
  ]
}
```

WKLD.03 – Use ephemeral secrets or a secrets-management service

Application secrets consist largely of credentials, such as key pairs, access tokens, digital certificates, and sign-in credentials. The application uses these secrets to gain access to other services it depends upon, such as a database. To help protect these secrets, we recommend they are either ephemeral (generated at the time of request and short-lived, such as with IAM roles) or retrieved from a secrets-management

service. This prevents accidental exposure through less secure mechanisms, such as persisting in static configuration files. This also makes it easier to promote application code from development to production environments.

For a secrets-management service, we recommend using a combination of Parameter Store, a capability of AWS Systems Manager, and AWS Secrets Manager:

- Use Parameter Store to manage secrets and other parameters that are individual key-value pairs, string-based, short in overall length, and accessed frequently. You use an AWS Key Management Service (AWS KMS) key to encrypt the secret. There is no charge to store parameters in the standard tier of Parameter Store. For more information about parameter tiers, see [Managing parameter tiers](#) (Systems Manager documentation).
- Use Secrets Manager to store secrets that are in document form (such as multiple, related key-value pairs), are larger than 4 KB (such as digital certificates), or would benefit from automated rotation.

You can use Parameter Store APIs to retrieve secrets stored in Secrets Manager. This allows you to standardize the code in your application when using a combination of both services.

To manage secrets in Parameter Store

1. [Create a symmetric AWS KMS key](#) (AWS KMS documentation).
2. [Create a SecureString parameter](#) (Systems Manager documentation). Secrets in Parameter Store use the `SecureString` data type.
3. In your application, retrieve a parameter from Parameter Store by using the AWS SDK for your programming language. For an example in Java, see [GetParameter.java](#) (AWS Code Sample Catalog).

To manage secrets in Secrets Manager

1. [Create a secret](#) (Secrets Manager documentation).
2. [Retrieve secrets from AWS Secrets Manager in code](#) (Secrets Manager documentation).

It is important to read [Use AWS Secrets Manager client-side caching libraries to improve the availability and latency of using your secrets](#) (AWS blog post). Using client-side SDKs, which already have best practices implemented, should accelerate and simplify the use and integration of Secrets Manager.

WKLD.04 – Prevent application secrets from being exposed

During local development, application secrets can be stored in local configuration or code files and accidentally checked-in to source code repositories. Unsecured repositories hosted on public service providers can be subject to unauthorized access and subsequent discovery of these secrets. Use available tools to prevent secrets from being checked in. Incorporate checks for exposed secrets as part of your manual code review processes.

Some common tools that can prevent application secrets from being checked-in to source code repositories are:

- [Gitleaks](#) (GitHub repository)
- [Whispers](#) (GitHub repository)
- [detect-secrets](#) (GitHub repository)

- [git-secrets](#) (GitHub repository)
- [TruffleHog](#) (GitHub repository)

WKLD.05 – Detect and remediate exposed secrets

In [WKLD.03 – Use ephemeral secrets or a secrets-management service \(p. 11\)](#) and [WKLD.04 – Prevent application secrets from being exposed \(p. 12\)](#), you put measures in place to protect secrets. In this control, you deploy a solution that can detect if secrets have bypassed these prevention measures, and you can remediate accordingly.

Amazon CodeGuru Reviewer detects application secrets in source code and provides a mechanism to remediate and publish the detected secrets in Secrets Manager. Application code for retrieving the secret from Secrets Manager is also provided. Conduct a cost-benefit analysis to determine if this solution is right for your business. As an alternative, some of the open-source solutions in [WKLD.04 – Prevent application secrets from being exposed \(p. 12\)](#) provide detection capability for existing secrets.

To set up CodeGuru Reviewer integration with Secrets Manager

- [Use CodeGuru Reviewer to identify hardcoded secrets and AWS Secrets Manager to secure them](#) (AWS blog post and guided walkthrough).

WKLD.06 – Use Systems Manager instead of SSH or RDP

Public subnets, which have a default route pointing to an internet gateway, are inherently a greater security risk than *private subnets*, those with no route to the internet. You can run EC2 instances in private subnets and use the Session Manager capability of AWS Systems Manager to remotely access the instances through either the AWS Command Line Interface (AWS CLI) or AWS Management Console. You can then use the AWS CLI or console to start a session that connects into the instance through a secure tunnel, preventing the need to manage additional credentials used for Secure Shell (SSH) or Windows remote desktop protocol (RDP).

Use Session Manager instead of running EC2 instances in public subnets, running jump boxes, or running bastion hosts.

To set up Session Manager

1. Make sure the EC2 instance is using the latest operating system Amazon Machine Images (AMIs), such as Amazon Linux 2 or Ubuntu. The AWS Systems Manager Agent (SSM Agent) is pre-installed on the AMI.
2. Make sure the instance has connectivity, either through an internet gateway or through VPC endpoints, to these addresses (replacing **<region>** with the appropriate AWS Region):
 - a. `Ec2messages.<region>.amazonaws.com`
 - b. `ssm.<region>.amazonaws.com`
 - c. `ssmmessages.<region>.amazonaws.com`
3. Attach the AWS managed policy `AmazonSSMManagedInstanceCore` to the IAM role that is associated to your instances.

For more information, see [Setting up Session Manager](#) (Systems Manager documentation).

To start a session

- [Start a session](#) (Systems Manager documentation).

WKLD.07 – Log data events for S3 buckets with sensitive data

By default, AWS CloudTrail captures management events, events that create, modify, or delete resources in your account. These management events do not capture read or write operations to individual objects in Amazon Simple Storage Service buckets. During a security event, it is important to capture unauthorized data access or use at an individual record or object level. Use CloudTrail to log data events for any S3 buckets that store sensitive or business-critical data, for detection and auditing purposes.

Note

Additional charges apply for logging data events. For more information, see [AWS CloudTrail pricing](#).

To log data events for trails

1. Sign in to the AWS Management Console and open the CloudTrail console at <https://console.aws.amazon.com/cloudtrail/>
2. In the navigation pane, choose **Trails**, and then choose a trail name.
3. In **General details**, choose Edit to change the following settings. You cannot change the name of a trail.
 - a. In **Data events**, choose **Edit**.
 - b. For **Data event source**, choose **S3**.
 - c. For **All current and future S3 buckets**, clear **Read** and **Write**.
 - d. In Individual bucket selection, browse for the bucket on which to log data events. You can select multiple buckets in this window. Choose **Add bucket** to log data events for more buckets. Choose to log **Read** events, such as GetObject, **Write** events, such as PutObject, or both.
 - e. Choose **Update trail**.

WKLD.08 – Encrypt Amazon EBS volumes

Enforce encryption of Amazon Elastic Block Store (Amazon EBS) volumes as the default behavior in your AWS account. Encrypted volumes have the same input/output operations per second (IOPS) performance as unencrypted volumes with a minimal effect on latency. This prevents rebuilding volumes at a later date for compliance or other reasons. For more information, see [Must-know best practices for Amazon EBS encryption](#) (AWS blog post).

To encrypt Amazon EBS volumes

- [Enable encryption by default](#) (Amazon EC2 documentation).

WKLD.09 – Encrypt Amazon RDS databases

Similar to [WKLD.08 – Encrypt Amazon EBS volumes \(p. 14\)](#), enable encryption of Amazon Relational Database Service (Amazon RDS) databases. This encryption is performed at the underlying volume level

and has the same IOPS performance as unencrypted volumes with a minimal effect on latency. For more information, see [Overview of encrypting Amazon RDS resources](#) (Amazon RDS documentation).

To encrypt an RDS database instance

- [Encrypt a database instance](#) (Amazon RDS documentation).

WKLD.10 – Deploy private resources into private subnets

Deploy resources that don't require direct internet access, such as EC2 instances, databases, queues, caching, or other infrastructure, into a VPC private subnet. Private subnets don't have a route declared in their route table to an attached internet gateway and cannot receive internet traffic. Traffic originating from a private subnet that is destined for the internet must undergo network address translation (NAT) through either a managed AWS NAT Gateway or an EC2 instance running NAT processes in a public subnet. For more information about network isolation, see [Infrastructure security in Amazon VPC](#) (Amazon VPC documentation).

Use the following practices when creating private resources and subnets:

- When creating a private subnet, disable **auto-assign public IPv4 address**.
- When creating private EC2 instances, disable **Auto-assign Public IP**. This prevents a public IP from being assigned if the instance is unintentionally deployed into a public subnet via misconfiguration.

You specify the subnet for a resource as part of its configuration, when required. You can deploy a VPC that follows best practices using the [Modular and Scalable VPC Architecture Quick Start](#) (AWS Quick Starts).

WKLD.11 – Restrict network access by using security groups

Use security groups to control traffic to EC2 instances, RDS databases, and other supported resources. Security groups act as a virtual firewall that can be applied to any group of related resources in order to consistently define rules for allowing inbound and outbound traffic. In addition to rules based on IP addresses and ports, security groups support rules to allow traffic from resources associated to other security groups. For example, a database security group can have rules to allow only traffic from an application server security group.

By default, security groups allow all outbound traffic but don't allow inbound traffic. The outbound traffic rule can be removed, or you can configure additional rules added to restrict outbound traffic and allow inbound traffic. If the security group has no outbound rules, no outbound traffic originating from your instance is allowed. For more information, see [Control traffic to resources using security groups](#) (Amazon VPC documentation).

In the following example, there are three security groups that control traffic from an Application Load Balancer to EC2 instances that connect to an Amazon RDS for MySQL database.

Security group	Inbound rules	Outbound rules
Application Load Balancer security group	Description: Allow HTTPS traffic from anywhere	Description: Allow all traffic to anywhere

Security group	Inbound rules	Outbound rules
	Type: HTTPS Source: Anywhere-IPv4 (0.0.0.0/0)	Type: All traffic Destination: Anywhere-IPv4 (0.0.0.0/0)
EC2 instance security group	Description: Allow HTTP traffic from the Application Load Balancer Type: HTTP Source: Application Load Balancer security group	Description: Allow all traffic to anywhere Type: All traffic Destination: Anywhere-IPv4 (0.0.0.0/0)
RDS database security group	Description: Allow MySQL traffic from EC2 instance Type: MySQL Source: EC2 instance security group	No outbound rules

WKLD.12 – Use VPC endpoints to access supported services

In VPCs, resources that need to access AWS or other external services require either a route to the internet (0.0.0.0/0) or to the public IP address of the target service. Use VPC endpoints to enable a private IP route from your VPC to supported AWS or other services, preventing the need to use an internet gateway, NAT device, virtual private network (VPN) connection, or AWS Direct Connect connection.

VPC endpoints support attaching policies and security groups to further control access to a service. For example, you can write a VPC endpoint policy for Amazon DynamoDB to allow only item-level actions and prevent table-level actions for all resources in the VPC, regardless of their own permission policy. You can also write an S3 bucket policy to allow only requests originating from a specific VPC endpoint, denying all other external access. A VPC endpoint can also have a security group rule that, for example, restricts access to only EC2 instances that are associated to an application-specific security group, such as the business-logic tier of a web application.

There are different kinds of VPC endpoints. You access most services by using a VPC interface endpoint. DynamoDB is accessed using a gateway endpoint. Amazon S3 supports both interface and gateway endpoints. Gateway endpoints are recommended for workloads contained within a single AWS account and Region, and come at no additional charge. Interface endpoints are recommended if more extensible access is required, such as to an S3 bucket from other VPCs, from on-premises networks, or from different AWS Regions. Interface endpoints incur an hourly uptime charge and a per-GB data-processing charge, both of which are lower than the respective charges for sending the data to 0.0.0.0/0 through AWS NAT Gateway.

See the following resources for additional information about using VPC endpoints:

- For more information about selecting between gateway and interface endpoints for Amazon S3, see [Choosing Your VPC Endpoint Strategy for Amazon S3](#) (AWS blog post).
- [Create an interface endpoint](#) (Amazon VPC documentation).

- [Create a gateway endpoint](#) (Amazon VPC documentation).
- For example S3 bucket policies that restrict access to a specific VPC or VPC endpoint, see [Restricting access to a specific VPC](#) (Amazon S3 documentation).
- For example DynamoDB endpoint policies that restrict actions, see [Endpoint policies for DynamoDB](#) (Amazon VPC documentation).

WKLD.13 – Require HTTPS for all public web endpoints

Require HTTPS to provide additional credibility to your web endpoints, allow your endpoints to use certificates to prove their identity, and confirm that all traffic between your endpoint and connected clients is encrypted. For public websites, this provides the additional benefit of higher search engine ranking.

Many AWS services provide public web endpoints for your resources, such as AWS Elastic Beanstalk, Amazon CloudFront, Amazon API Gateway, Elastic Load Balancing, and AWS Amplify. For instructions about how require HTTPS for each of these services, see the following:

- [Elastic Beanstalk](#) (Elastic Beanstalk documentation)
- [CloudFront](#) (CloudFront documentation)
- [Application Load Balancer](#) (AWS Knowledge Center)
- [Classic Load Balancer](#) (AWS Knowledge Center)
- [Amplify](#) (Amplify documentation)

Static websites hosted on Amazon S3 do not support HTTPS. To require HTTPS for these websites, you can use CloudFront. Public access to S3 buckets that are serving content through CloudFront is not required.

To use CloudFront to serve a static website hosted on Amazon S3

1. [Use CloudFront to serve a static website hosted on Amazon S3](#) (AWS Knowledge Center).
2. If you are configuring access to a public S3 bucket, [require HTTPS between viewers and CloudFront](#) (CloudFront documentation).

If you are configuring access to a private S3 bucket, [restrict access to Amazon S3 content by using an origin access identity](#) (CloudFront documentation).

In addition, configure HTTPS endpoints to require modern Transport Layer Security (TLS) protocols and ciphers, unless compatibility with older protocols is needed. For example, use the `ELBSecurityPolicy-FS-1-2-Res-2020-10` or the most recent policy available for Application Load Balancer HTTPS listeners, instead of the default `ELBSecurityPolicy-2016-08`. The most current policies require TLS 1.2 at minimum, forward secrecy, and strong ciphers that are compatible with modern web browsers.

For more information about the available security policies for HTTPS public endpoints, see:

- [Predefined SSL security policies for Classic Load Balancers](#) (Elastic Load Balancing documentation)
- [Security policies for your Application Load Balancer](#) (Elastic Load Balancing documentation)
- [Supported protocols and ciphers between viewers and CloudFront](#) (CloudFront documentation)

WKLD.14 – Use edge-protection services for public endpoints

Rather than serve traffic direct from compute services such as EC2 instances or containers, use an edge-protection service. This provides an additional layer of security between incoming traffic from the internet and your resources that serve that traffic. These services can filter unwanted traffic, enforce encryption, and apply routing or other rules, such as load balancing, before traffic reaches your internal resources.

AWS services that can provide public endpoint protection include the AWS WAF, CloudFront, Elastic Load Balancing, API Gateway, and Amplify Hosting. Run VPC-based services, such as Elastic Load Balancing, in a public subnet as a proxy to web service resources running in a private subnet.

CloudFront, API Gateway, and Amazon Route 53 provide protection from Layer 3 and 4 distributed denial of service (DDoS) attacks at no charge, and AWS WAF can protect against Layer 7 attacks.

Instructions for getting started with each of these services can be found here:

- [Getting Started with AWS WAF](#) (AWS website)
- [Getting started with Amazon CloudFront](#) (CloudFront documentation)
- [Getting started with Elastic Load Balancing](#) (Elastic Load Balancing documentation)
- [Getting started with API Gateway](#) (API Gateway documentation)
- [Getting started with Amplify Hosting](#) (Amplify documentation)

WKLD.15 – Define security controls in templates and deploy them by using CI/CD practices

Infrastructure as code (IaC) is the practice of defining all of your AWS service resources and configurations in templates and code that you deploy by using continuous integration and continuous delivery (CI/CD) pipelines, the same pipelines used to deploy software applications. IaC services, such as AWS CloudFormation, support both IAM identity-based and resource-based policies and support AWS security services, such as Amazon GuardDuty, AWS WAF, and Amazon VPC. Capture these artifacts as IaC templates, commit the templates to a source code repository, and then deploy them by using CI/CD pipelines.

Unless required otherwise, commit application permission policies with application code in the same repository, and manage general resource policies and security service configurations in separate code repositories and deployment pipelines.

For more information about getting started with IaC on AWS, see the [AWS Cloud Development Kit \(AWS CDK\) documentation](#).

Contributors

Contributors to this document include:

- Jay Michael, Principal Solutions Architect
- Cole Calistra, Principal Solutions Architect
- Justin Plock, Principal Solutions Architect
- Faisal Farooq, Solutions Architect
- Michael Nguyen, Sr. Solutions Architect
- Ritik Khatwani, Sr. Solutions Architect
- Paul Hawkins, Principal, Office of the Chief Information Security Officer (CISO)

A special thank you to the following people who also helped with guidance and review:

- Robert Put
- Mike Sullivan
- Bob Lee III

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
Amazon S3 bucket settings (p. 20)	We updated the ACCT.08 – Prevent public access to private S3 buckets section to reflect that Amazon S3 buckets created after April 28, 2023 have the Block Public Access setting enabled by default.	May 18, 2023
IAM security best practices (p. 20)	We updated this guide for alignment with the latest AWS Identity and Access Management (IAM) best practices. For more information, see Security best practices in the IAM documentation.	February 1, 2023
IAM roles (p. 20)	We provided additional links to AWS service documentation in the WKLD.01 – Use IAM roles for compute environment permissions section.	September 22, 2022
Password policy (p. 20)	We updated the recommendations for strong passwords to use the latest guidance from the Center for Internet Security (CIS).	May 10, 2022
Initial publication (p. 20)	—	April 13, 2022

AWS Prescriptive Guidance glossary

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

Security terms

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.

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.

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.

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.

client-side encryption

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

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.

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 in transit

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

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 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.

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.

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*.

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.

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 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 VPC documentation.

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.

fine-grained access control (FGAC)

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

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.

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.

identity-based policy

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

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.

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.

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.

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.

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.

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.

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 \(p. 23\)](#), which provides more granular and enhanced access control.

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.

policy

An object that can define permissions (see [identity-based policy \(p. 23\)](#)), specify access conditions (see [resource-based policy \(p. 24\)](#)), or define the maximum permissions for all accounts in an organization in AWS Organizations (see [service control policy \(p. 25\)](#)).

preventative control

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

principal

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

Privacy by Design

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

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.

ransomware

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

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.

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*.

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.

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 three primary types of security controls: [preventative \(p. 24\)](#), [detective \(p. 22\)](#), and [responsive \(p. 24\)](#).

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.

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.

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](#).

symmetric encryption

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

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.

vulnerability

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

workload

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

zero-day exploit

An attack, typically malware, that takes advantage of a [zero-day vulnerability \(p. 25\)](#).

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.