AWS Prescriptive Guidance

Strategy for securing semiconductor development environments on AWS

AWS Prescriptive Guidance: Strategy for securing semiconductor development environments on AWS

Table of Contents

Strategy for securing semiconductor development environments on AWS

Mike Virgilio, Allan Carter, and Nikhil Marrapu, Amazon Web Services (AWS)

June 2023 (document history (p. 16))

This document provides strategic guidance to help you secure and meet compliance requirements for semiconductor workloads operating in the AWS Cloud. It includes sample architectures and overviews of AWS services and features that can you can use to implement this guidance and help protect sensitive data from security risks.

Semiconductor industry overview

According to the Semiconductor Industry Association, the semiconductor business is a USD573.44 billion global industry. Security is a top priority for the industry in order to safeguard intellectual property (IP). Semiconductor companies require development environments so their engineers can design chips, electronic systems, circuit boards, and other products. A *secure development environment* must strictly control who can access the IP within it.

Companies develop their own IP, but they commonly also use IP from third-party vendors, such as processor cores, standard interfaces, process design kits (PDKs) from semiconductor foundries, and licensed tools from electronic design automation (EDA) companies. The highly collaborative nature of the development process means internal engineers and engineers from those third-party companies need access to the development environment. A critical security requirement is to protect against unauthorized data exfiltration from the secure development environment.

Semiconductor companies require the latest technical capabilities

Semiconductor companies operate within a competitive industry, where speed-to-market and innovation are essential for success. As chip design and fabrication requirements become more intricate, semiconductor companies require access to the latest technologies to meet and exceed the industry's demands. The exponential growth in compute and storage requirements can be met by the scalability and capacity of the AWS Cloud. With a comprehensive infrastructure and robust set of compute, network, and storage solutions, AWS empowers semiconductor companies to utilize cutting-edge technologies, such as machine learning, high-performance computing, and automation. Use of these technologies can accelerate research and development efforts, optimize the manufacturing processes, and provide access to the latest technology. Valuable IP is a compelling target for sophisticated attacks, making security the top priority for a secure development environment.

Using economies of scale, elasticity, and automation

AWS provides companies with economy of scale, resource elasticity, and automation capabilities that are essential to success. Because AWS has partnered with hundreds of thousands of companies, massive

economies of scale can be achieved, and this translates to lower costs for all. AWS infrastructure elasticity allows companies to easily scale up to satisfy the most demanding workloads and then scale down to optimize costs. In addition, AWS automation capabilities help companies create repeatable processes that minimize undifferentiated, manual tasks. AWS offers a wide range of security services and features to help semiconductor companies secure their workloads through strong security controls, including network segmentation, data encryption, and regulatory compliance. By building in the AWS Cloud, semiconductor companies can focus on innovation and growth, while also ensuring that their data and operations are resilient against potential security risks.

Achieving security and compliance for semiconductor development environments on AWS

AWS has developed best practice guidance to implement security controls and published reference architectures to address semiconductor industry needs. This section discusses how to use the AWS recommended designs and reference architectures to help achieve security and compliance for your mission-critical workloads on AWS.

Reducing compliance efforts with AWS

The AWS shared responsibility model describes how responsibility for security and compliance is shared between AWS and the customer. AWS is responsible for security *of* the cloud, and the customer is responsible for security *in* the cloud. This can help companies reduce the effort necessary to achieve compliance with corporate and regulatory requirements by placing the responsibility for cloud infrastructure on AWS.

The following AWS services can help semiconductor companies demonstrate compliance with corporate and regulatory requirements:

- AWS Artifact provides downloadable compliance reports for various compliance frameworks, including International Organization for Standardization (ISO), National Institute of Standards and Technology (NIST), and Federal Risk and Authorization Management Program (FedRAMP). You can combine AWS Artifact reports with corporate assessment of cloud resources to demonstrate compliance to auditors and help reduce the time and effort required to become compliant with regulations such as United States International Traffic in Arms Regulations (ITAR).
- AWS Audit Manager can map your compliance requirements to AWS usage data by using prebuilt and custom frameworks and automated evidence collection.

By using these services and features, companies can achieve compliance with corporate and regulatory requirements more efficiently and effectively. For more information about whether an AWS service is in scope of AWS assurance programs, see AWS services in scope by compliance program.

Using provided reference architectures

AWS develops prescriptive guidance and best practices based on thousands of deployments across various industries. These recommendations are included within the AWS Well-Architected Framework, AWS Cloud Adoption Framework (AWS CAF), and AWS Security Reference Architecture (AWS SRA).

When architecting and designing your secure development environment, AWS provides semiconductor and electronics reference architectures that are based on the aforementioned frameworks. These reference architectures are designed to protect data and workloads.

You can use the AWS Security Maturity Model to guide you through the backlog of security controls in a phased approach.

By utilizing these frameworks, models, and reference architectures, you can establish a robust security posture in the cloud and help protect critical assets.

Securing semiconductor data on AWS

With AWS, you manage the privacy controls of your data, and you control how your data is used, who has access to it, and how it is encrypted. These capabilities are underpinned with a flexible and secure cloud computing environment. This section reviews AWS capabilities to help semiconductor companies implement data access controls and address data residency requirements.

This section contains the following topics:

Preventing unauthorized access and data exfiltration

According to 2022 Cost of a Data Breach Report (Ponemon Institute report), the average cost of a data breach in 2022 was USD4.3 million. For semiconductor companies, protecting intellectual property (IP) is critical. Loss of IP due to unauthorized access can result in financial loss, reputational damage, or even regulatory consequences. These potential consequences make controlling access to the data and the flow of data critical aspects of a well-architected design.

Key considerations to secure your data include:

- User authentication for access to the secure development environment
- User authorization for access to data within the secure development environment
- Logging all transfers into and out of the secure development environment
- Architecting secure data flows between environments
- Encryption of data in transit and rest
- Limiting and logging outbound network traffic

Configuring permissions

AWS Identity and Access Management (IAM) helps you securely manage access to your AWS resources by controlling who is authenticated and authorized to use them. By default, any action in AWS is implicitly denied unless it is explicitly allowed. You manage access in AWS by creating *policies*. You can use policies to define, at a granular level, which users can access which resources and what actions they can perform on those resources. An AWS best practice is to apply least-privilege permissions, which means you grant users only the permissions they require to perform their tasks. For more information, see the following in the IAM documentation:

- Policies and permissions in IAM
- Policy evaluation logic
- Security best practices in IAM

Authenticating users

It's an AWS best practice to require human users to use federation with an identity provider to access AWS by using temporary credentials. The recommended service for centralizing your user workforce access is AWS IAM Identity Center. This service helps you securely create or connect your workforce identities and manage their access centrally across AWS accounts and applications. IAM Identity Center can federate with external identity providers (IdPs) by using SAML 2.0, Open ID Connect (OIDC), or OAuth 2.0 in order to provide seamless integration and user management. For more information, see Identity federation in AWS (AWS marketing) and Identity providers and federation (IAM documentation).

You can also authenticate and authorize users by using AWS Directory Service to manage users and groups that are defined in a directory, such as Active Directory. Within the secure development environment, you can use Linux file permissions to authorize and restrict data access within the virtual private cloud (VPC). Use VPC endpoints to provide access to AWS services without traversing the public internet. Use endpoint policies to restrict which AWS principals can use the endpoint, and use identity-based policies to restrict access to AWS services.

Transferring data

AWS provides several ways to migrate on-premises data to the cloud. It's common to initially store the data in Amazon Simple Storage Service (Amazon S3). Amazon S3 is a cloud-based object storage service that helps you store, protect, and retrieve any amount of data. It provides bandwidth of up to 25 Gbps when transferring data to or from an Amazon Elastic Compute Cloud (Amazon EC2) instance. It also offers cross-Region data replication and data tiering. Data stored in Amazon S3 can serve as a replication source. You can use it to create new file systems or to transfer data to EC2 instances. You can use Amazon S3 as the backend of an AWS managed, Portable Operating System Interface (POSIX)-compliant file system for semiconductor tools and flows.

Another AWS storage service is Amazon FSx, which provides file systems that support industry-standard connectivity protocols and offers high availability and replication across AWS Regions. Common choices for the semiconductor industry include Amazon FSx for NetApp ONTAP, Amazon FSx for Lustre, and Amazon FSx for OpenZFS. The scalable, high-performance file systems in Amazon FSx are well suited to storing data locally within the secure development environment.

AWS recommends that you define storage requirements for your semiconductor workloads on AWS first and then identify the appropriate data transfer mechanism. AWS recommends using AWS DataSync to transfer data from on premises to AWS. DataSync is an online data transfer and discovery service that helps you move files or object data to, from, and between AWS storage services. Depending on whether you are using self-managed storage systems or a storage provider such as NetApp, you can configure DataSync to accelerate moving and replicating data to your secure development environment over the internet or through AWS Direct Connect. DataSync can transfer your file system data and metadata, such as ownership, timestamps, and access permissions. If you're transferring files between FSx for ONTAP and NetApp ONTAP, AWS recommends using NetApp SnapMirror. Amazon FSx supports encryption at rest and in transit. Use AWS CloudTrail and other service-specific logging features to log all API calls and related data transfers. Centralize logs in a dedicated account, and apply granular access policies for immutable history.

AWS provides additional services to help control data flows, including application-aware network firewalls such as AWS Network Firewall, Amazon Route 53 Resolver DNS Firewall, AWS WAF, and web proxies. Control data flows within the environment by enforcing network segmentation with security groups, network access control lists, and VPC endpoints in Amazon Virtual Private Cloud (Amazon VPC), Network Firewall, transit gateway route tables, and service control policies (SCPs) in AWS Organizations. Centrally log all network traffic by using VPC Flow Logs and the available fields from versions 2–5 of VPC Flow Logs.

Encrypting data

Encrypt all data at rest by using AWS Key Management Service (AWS KMS) customer managed keys or AWS CloudHSM. Create and maintain granular key resource policies. For more information, see Creating an enterprise encryption strategy for data at rest.

Encrypt data in transit by enforcing a minimum of TLS 1.2 with an industry-standard 256-bit Advanced Encryption Standard (AES-256) cipher.

Managing outbound network traffic

If the secure development environment requires internet access, then all outbound internet traffic should be logged and restricted through a network-level enforcement point, such as through Network Firewall or Squid, which is an open-source proxy. VPC endpoints and the internet proxy help protect against unauthorized exfiltration of data by users. This is critical to allow access to data within the secure development environment and only within the VPC.

Finally, you can use Network Access Analyzer, a feature of Amazon VPC, to perform network segmentation validation and identify potential network paths that do not meet your specified requirements.

By layering of security controls, you can establish and enforce a robust data perimeter. For more information, see Building a Data Perimeter on AWS.

Meeting data residency requirements on AWS

The available partitions, AWS Regions, Availability Zones, and Local Zones allow companies to choose the best location for their data and workloads based on their unique requirements:

- A partition is a logical group of AWS Regions. AWS commercial Regions are in the `aws` partition, Regions in China are in the `aws-cn` partition, and AWS GovCloud (US) Regions are in the `aws-us-gov` partition.
- An AWS Region is a separate geographic area where AWS clusters data centers.
- Each AWS Region has multiple, isolated locations known as Availability Zones.
- A Local Zone is an extension of a Region that is geographically close to your users.

For more information about the currently available Regions, Availability Zones, and Local Zones, see AWS Global Infrastructure.

A partition provides data, network, and machine isolation from Regions in other partitions. AWS partitions create logical network isolation with separate credentialed access between Regions in the different partitions. Partitions include one or more Regions, but an AWS Region exists only within one partition; an AWS Region cannot be a part of two partitions.

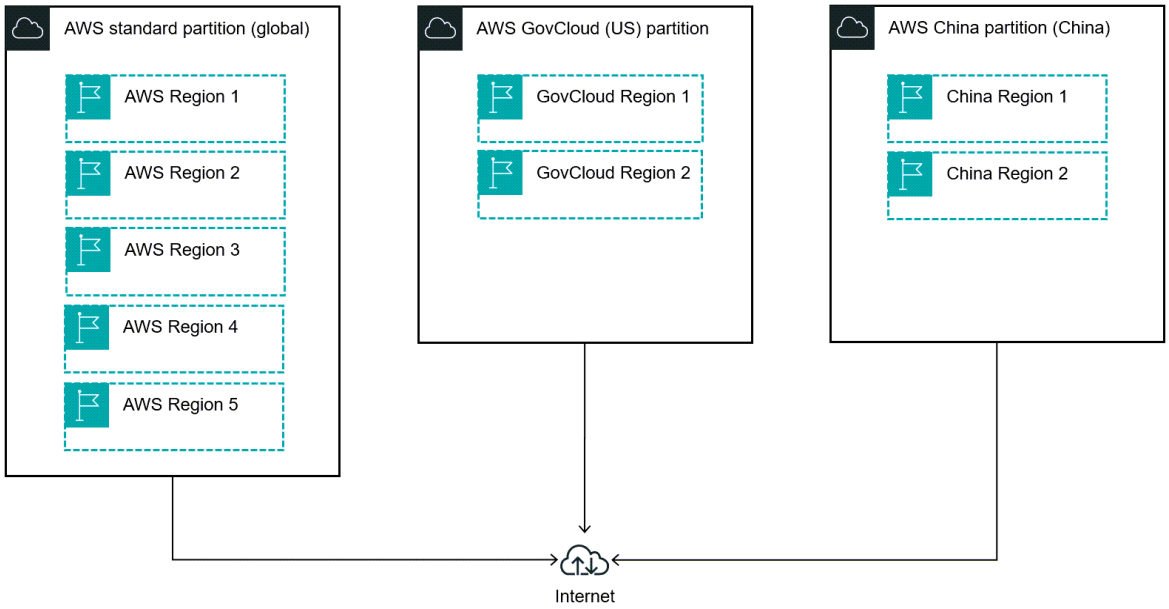

You can choose between partitions based on whether a United States government security classification is required. Workloads processing unclassified or official data can use both the AWS GovCloud (US) or standard partitions. AWS also offers additional partitions accredited to operate workloads at the Secret and Top-Secret US security classification levels, but these are out of scope for this guide. For more information about operating workloads at these classification levels, see Cloud Computing for US Defense and Cloud Computing for the US Intelligence Community.

We recommend deploying multi-Region workloads within a single partition to reduce any compliance, operational, and technical challenges. However, there are limited use cases, such as with AWS Direct Connect or Amazon CloudFront, where you can integrate services across multiple to meet specific objectives. For more information, contact your AWS Solutions Architect.

Providing a secure remote experience for engineers

Remote visualization is a key component in semiconductor design. Tool engineers use remote verification, and chip designers use it to submit jobs and view their status. This section presents a robust solution that can perform well over varying network conditions. It is designed be cloud-native so that it can seamlessly integrate with AWS security services.

This section contains the following topics:

Optimize the remote desktop experience

Designers typically use terminal-based SSH sessions or graphical remote desktops to submit and visualize workflows. A remote desktop offers GUI-driven interactive tools (such as layout, place, and route) for tool engineers and chip designers to submit jobs. AWS offers NICE DCV, which is a high-performance remote display protocol that provides a robust user interface for engineering and physical design teams. NICE DCV performs well over varying network conditions. The NICE DCV for Amazon Linux 2 Amazon Machine Image (AMI) is available on AWS Marketplace for no charge.

NICE DCV streams pixels and not geometries in order to help protect data privacy. In addition, NICE DCV uses TLS to secure pixels and end-user inputs.

Using a connection file, users can instantly connect to a NICE DCV session. However, note that the connection file parameters use the `password` and `proxypassword` fields without encryption. For more information, see Using a connection file. NICE DCV establishes a TLS connection between the server and client. A validation policy in the connection file determines how the client responds when a certificate can't be verified as trustworthy . For more information, see Set certificate validation policy.

Other on-premises commercial solutions that provide remote desktop functionality include NoMachine or OpenText Exceed TurboX.

With any remote desktop solution, the underlying infrastructure is powered by Amazon Elastic Compute Cloud (Amazon EC2). According to the shared responsibility model, your responsibility includes the following areas to help secure remote desktop instances:

- Controlling network access to your instances, such as by configuring your VPC and security groups. For more information, see Controlling network traffic.
- Managing the credentials used to connect to your instances.
- Managing the guest operating system and software deployed to the guest operating system, including updates and security patches. For more information, see Update management in Amazon EC2.
- Configuring the IAM roles that are attached to the instance and the permissions associated with those roles. For more information, see IAM roles for Amazon EC2.

Securing engineering collaboration with third parties

Collaboration with third parties is essential during development in order to debug tool problems, get help with IP integration for designs, and to bring in external contractors with specialized skills. It can be difficult to provide secure access to third parties from on-premises infrastructure. By using AWS infrastructure as code (IaC), you can create a copy of your primary secure development environment, called a *collaboration chamber*. To help prevent data exfiltration, tighten the security posture of a collaboration chamber by not allowing access to the internet. A collaboration chamber has accounts for the collaborators, and you can curate the data, tools, and infrastructure in the chamber to include only what is required for the collaboration. When the collaboration is done, delete the collaboration chamber in order to reduce costs and to remove any potential access to the data. The following diagram shows how different participants in the design and manufacturing process might use various types of collaboration chambers.

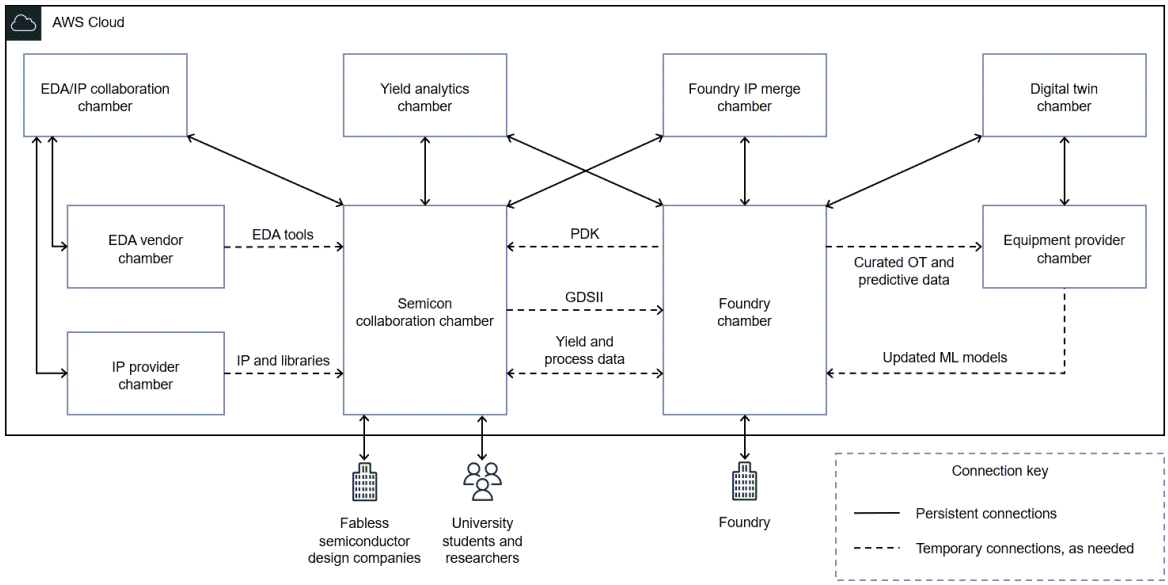

The following image is a reference architecture for a collaboration chamber. This architecture can be used as reference when designing and building a collaboration chamber on AWS. The AWS security, governance, and monitoring services in the diagram help secure the chamber in order to protect IP. For more information about these services, see AWS security services for semiconductor development environments (p. 12) in this guide.

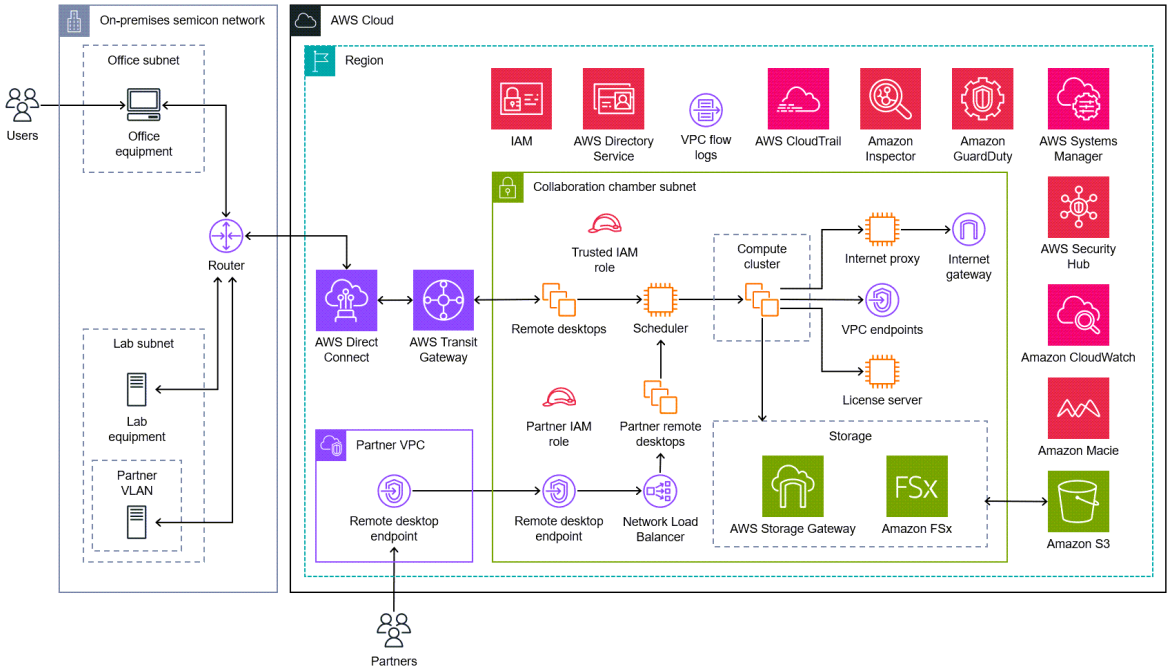
On-premises semicon network
Office subnet
Users
Office equipment
Router
Lab subnet
Lab equipment
Partner VLAN
AWS Cloud
Region
IAM
AWS Directory Service
VPC flow logs
AWS CloudTrail
Amazon Inspector
Amazon GuardDuty
AWS Systems Manager
Collaboration chamber subnet
Trusted IAM role
Compute cluster
Internet proxy
Internet gateway
AWS Security Hub
AWS Direct Connect
AWS Transit Gateway
Remote desktops
Scheduler
VPC endpoints
Amazon CloudWatch
License server
Partner IAM role
Partner remote desktops
Partner VPC
Storage
Amazon Macie
Remote desktop endpoint
Remote desktop endpoint
Network Load Balancer
AWS Storage Gateway
Amazon FSx
Amazon S3
Partners

AWS security services for semiconductor development environments

AWS has developed security services designed to protect your workloads and applications in the AWS Cloud. These services can help protect any type of workload operating in the cloud, not just semiconductor workloads. By using these AWS services, companies can establish strong preventative and detective controls, monitor their security posture in near real-time, and quickly remediate security risks and incidents as they arise. These services are capable of automatically scaling with the environment as it grows in order to maintain coverage and the established security posture. Also, while these services might focus on specific functionality, they support a common messaging bus known as Amazon EventBridge so that you can integrate the services and automatically response to security risks.

The following AWS services and features can help you manage access and policies, detect and respond to security risks and events, and implement monitoring and logging in your AWS environment:

- AWS CloudTrail helps you audit the governance, compliance, and operational risk of your AWS account.
- Amazon GuardDuty is a continuous security monitoring service that analyzes and processes logs to identify unexpected and potentially unauthorized activity in your AWS environment.
- Amazon Inspector is a vulnerability management service that continuously scans your AWS workloads for software vulnerabilities and unintended network exposure.
- AWS Security Hub provides a comprehensive view of your security state in AWS. It also helps you check your AWS environment against security industry standards and best practices.
- Amazon Security Lake automatically centralizes security data from cloud, on-premises, and custom sources into a purpose-built data lake stored in your AWS account. You can query and analyze this security data to discover trends and anomalies.
- Service control policies (SCPs) are a type of policy in AWS Organizations that helps you centrally manage the use of AWS services across multiple accounts.
- VPC Flow Logs is a feature of Amazon Virtual Private Cloud (Amazon VPC) that captures information about the IP traffic going to and from network interfaces in your VPC. You can use this data to troubleshoot and respond to incidents.

Securing semiconductor workloads and data against ransomware

Ransomware is a malicious software that is designed to block access to a computer system or data until a payment is made. Unfortunately, this type of attack is prevalent across the internet and can affect enterprises of any size. Aside from financial loss and reputational damage, semiconductor companies must prevent operational disruption because disruptions can cause large, cascading supply chain affects across the business ecosystem. It can also result in the loss of valuable intellectual property.

To help protect against ransomware, semiconductor companies must be vigilant to secure all aspects of their on-premises and cloud environments, especially systems that store, process, or transmit data. AWS provides many architecture designs and controls that can help protect against these types of incidents. In addition to the controls mentioned in this guide, see the following AWS resources about protecting cloud workloads from ransomware:

- AWS Cloud Security - Protecting against ransomware (AWS Cloud Security)
- AWS Blueprint for Ransomware Defense (AWS Security blog post)
- Protecting your AWS environment from ransomware (AWS webinar)
- eBook: Protecting your AWS environment from ransomware (AWS eBook)
- The anatomy of ransomware event targeting data residing in Amazon S3 (AWS Security blog post)
- Ransomware Risk Management on AWS Using the NIST Cyber Security Framework (CSF) (AWS whitepaper)

The AWS Cloud differentiator for the semiconductor industry

The business case for migrating from on-premises environments to AWS can be highly compelling for semiconductor companies. On-premises infrastructure can be expensive to manage, requires significant upfront investments, requires a significant amount of effort to maintain and stay current with the latest technology, and requires sourcing security tooling from a wide range of independent vendors. With AWS, you pay only for the resources you use, which can help you optimize costs for workloads. AWS also provides instant access to the latest high-performance computing (HPC) and cloud storage services and features. As an example, Amazon FSx for NetApp ONTAP is a commonly used storage service. You can use it to seamlessly migrate data from on-premises NetApp storage to the AWS Cloud, and it provides network isolation, resource-level permissions, identity-based authentication, encryption, logging, and auditing.

AWS can also improve the efficiency of provisioning and managing security solutions for semiconductor companies. Using the advanced security features and services in AWS, you can implement a robust security posture that helps protect your data and workloads against security risks. By using AWS Organizations and integrated security services, you can automatically deploy security controls across all AWS accounts within the organization. Automation capabilities help you automatically respond to risks and streamline your security operations, reducing the time and resources required to manage comprehensive security architectures. For example, AWS provides tools and services for automated security management, continuous compliance monitoring, and risk detection and response across the entire environment. By using these solutions, semiconductor companies can improve their security posture while also reducing costs and increasing efficiency. Overall, migrating to AWS can help semiconductor companies realize significant benefits in terms of scalability, flexibility, and security.

Conclusion

AWS provides an extensive and secure platform that supports the needs of semiconductor companies. By using the latest technologies and by taking advantage of the economy of scale, elastic scaling capabilities, and a wide range of security services, semiconductor companies can transform their business in the AWS Cloud. These benefits of the AWS Cloud solve complex challenges while optimizing for innovation and continued growth. AWS global infrastructure is prepared to support the most sophisticated requirements, such as data residency and security classification requirements. AWS can help you reduce costs, improve agility, innovate faster, and secure intellectual property in the cloud. For help with configuring semiconductor development environments in the AWS Cloud, contact your AWS account team or use the Contact AWS form. For additional AWS information and resources for the semiconductor industry, see Semiconductor & Hi-Tech Electronics.

Document history

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

Change	Description	Date
Initial publication (p. 16)	—	June 20, 2023

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.

Management and governance terms

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 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.

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.

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*.

RASCI matrix

: See RACI matrix (p. 17).

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.

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.

Networking terms

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.

endpoint

See service endpoint (p. 18).

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.

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.

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 Managing AWS Regions in *AWS General Reference*.

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*.

subnet

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

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.

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.

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. 21), 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. 20)), specify access conditions (see resource-based policy (p. 22)), or define the maximum permissions for all accounts in an organization in AWS Organizations (see service control policy (p. 22)).

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. 21), detective (p. 20), and responsive (p. 22).

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. 23).

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.

Storage and backup terms

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.

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.

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 (p. 23). For more information, see Disaster Recovery of Workloads on AWS: Recovery in the Cloud in the AWS Well-Architected Framework.

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.

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.

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.

warm data

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