AWS Compliance Guide to NYDFS Cybersecurity Regulation

Second Amendment to NYDFS Cybersecurity Regulation
Part 500 of Title 23

July 2024

Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents current AWS product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers or licensors. AWS products or services are provided “as is” without warranties, representations, or conditions of any kind, whether express or implied. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

Additionally, this document does not constitute legal advice and should not be relied on as legal advice. AWS encourages its customers to obtain appropriate advice on their implementation of privacy and data protection environments, and more generally, applicable laws relevant to their business.

Contents

Abstract

Entities in the State of New York operating under, or required to operate under, a license, registration, charter, certificate, permit, accreditation, or similar authorization under the Banking Law, the Insurance Law, or the Financial Services Law need to comply with the requirements in the Second Amendment to the NYDFS Cybersecurity Regulation (Part 500 of Title 23).

The entities in scope for the NYDFS Cybersecurity Regulation, known as *Covered Entities*, can use this guide when they evaluate their compliance to the requirements in the regulation. This guide describes the roles that AWS and Covered Entities play in managing and securing the cloud environment, describes the AWS Shared Responsibility Model, and provides additional resources that Covered Entities can use when they design and architect their AWS environments to meet their regulatory objectives under the NYDFS Cybersecurity Regulation.

Introduction

Around the world, financial institutions (FIs) use Amazon Web Services (AWS) to modernize and automate their core applications, including mobile applications, regulatory reporting, and market analysis. Through continuous innovation, AWS makes strong security available to FIs globally, along with a deep set of services and features, industry expertise, and the AWS Partner Network. AWS empowers FIs to modernize their technology infrastructure, meet rapidly changing customer behaviors and expectations, and drive business growth. AWS offers IT services in categories from compute, storage, database, and networking to artificial intelligence and machine learning.

Effective March 1, 2017, the New York State Superintendent of Financial Services promulgated 23 NYCRR Part 500, a regulation with cybersecurity requirements for financial services companies (*Cybersecurity Regulation* or *Part 500*). The entities required to comply with the Cybersecurity Regulation include, but are not limited to, partnerships, corporations, branches, agencies, and associations operating under, or required to operate under, a license, registration, charter, certificate, permit, accreditation, or similar authorization under the Banking Law, the Insurance Law, or the Financial Services Law (*Covered Entities*).

This guide is a resource to help Covered Entities understand technical and operational requirements for the use of AWS services. This guide includes a description of the tools and security features offered by AWS that Covered Entities can use to assist them with compliance with requirements in the Second Amendment to NYDFS Cybersecurity Regulation Part 500 of Title 23 (*NYDFS Cybersecurity Regulation*).

This guide does not undertake a full analysis of the NYDFS Cybersecurity Regulation. The sections in the following list provide information on AWS services, features, and resources that can help Covered Entities support their regulatory objectives under the NYDFS Cybersecurity Regulation.

- **Security and shared responsibility**: It is important that Covered Entities understand the AWS Shared Responsibility Model before evaluating the specific technical and operational requirements outlined in the NYDFS Cybersecurity Regulation. The AWS Shared Responsibility Model is fundamental to understanding the respective roles of the Covered Entity and AWS with respect to security and information access.
- **AWS Compliance Programs**: AWS has obtained certifications and third-party attestations for a variety of industry-specific and general workloads. AWS has also developed Compliance Programs to make these resources available to customers. Customers can take advantage of AWS Compliance Programs to help satisfy their regulatory objectives.

- **AWS Global Cloud Infrastructure**: The AWS Global Cloud Infrastructure comprises AWS Regions and Availability Zones. The AWS Global Cloud Infrastructure offers AWS customers a way to design and operate applications and databases, making them more available, fault tolerant, and scalable than traditional on-premises environments. AWS customers can use the AWS Global Cloud Infrastructure to help them design an AWS environment consistent with their business and regulatory objectives.
- **Appendix: Considerations on the Second Amendment to NYDFS Cybersecurity Regulation** Describes considerations for Covered Entities that use AWS and describes how Covered Entities can use AWS services and tools to support their regulatory objectives under the NYDFS Cybersecurity Regulation.

This document is provided for informational purposes only; it is not legal or compliance advice and should not be relied on as legal or compliance advice. Covered Entities are responsible for making their own independent assessments and should obtain appropriate advice from their own legal and compliance advisors regarding compliance with applicable regulations.

Security and the AWS Shared Responsibility Model

Cloud security is a shared responsibility and Covered Entities need to understand the AWS Shared Responsibility Model before reviewing their requirements under the NYDFS Cybersecurity Regulation. AWS manages security *of* the cloud by maintaining the AWS Cloud Infrastructure aligned with global and regional regulatory requirements and best practices. Security *in* the cloud is the responsibility of the Covered Entity. Namely, Covered Entities retain control of the security programs that they choose to implement to protect their content, applications, systems, and networks. Covered Entities assume responsibility and management of the guest operating system (including updates and security patches), and other associated application software in addition to the configuration of the security group firewall provided by AWS.

Covered Entities should carefully consider the services they choose because their responsibilities vary depending on the services used, the integration of those services into their IT environment, and applicable laws and regulations. As shown in Figure 1, this differentiation of responsibility between AWS and AWS customers is referred to as Security *of* the cloud versus Security *in* the cloud.

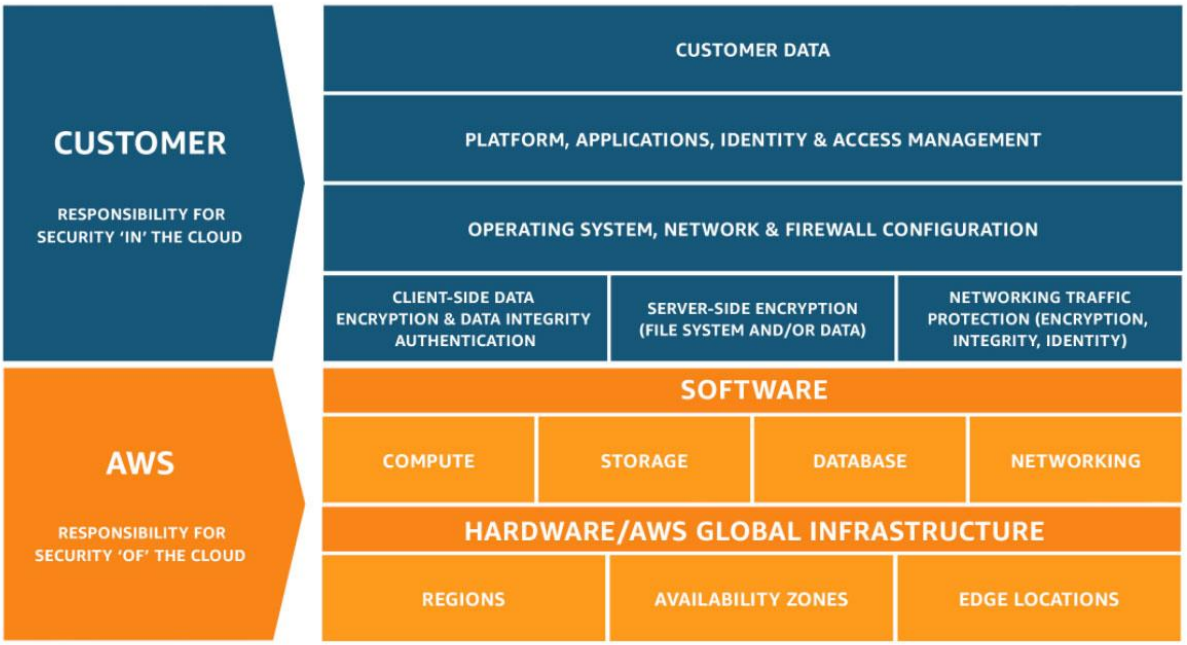

Figure 1 – The AWS Shared Responsibility Model

AWS responsibility - Security of the cloud: AWS is responsible for protecting the infrastructure that runs the AWS services. This infrastructure is composed of the hardware, software, networking, and facilities that run AWS services.

Customer responsibility - Security in the cloud: Customer responsibility is determined by the AWS services that a customer selects. This determines the amount of configuration work the

customer must perform as part of their security responsibilities. For example, a service such as Amazon Elastic Compute Cloud (Amazon EC2) requires the customer to perform all of the necessary security configuration and management tasks. Customers that deploy an EC2 instance are responsible for management of the guest operating system (including updates and security patches), application software or utilities installed by the customer on the EC2 instances, and the configuration of the firewall provided by AWS (called a security group) on each EC2 instance.

For abstracted services, such as Amazon Simple Storage Service (Amazon S3) and Amazon DynamoDB, AWS operates the infrastructure layer, the operating system, and environments, and customers access the endpoints to store and retrieve data. Customers are responsible for managing their data (including encryption options), classifying their assets, and using AWS Identity and Access Management (IAM) tools to apply the appropriate permissions.

It is important to note that when using AWS services, customers maintain control over their content and are responsible for managing critical content security requirements, including:

- The content that they choose to store on AWS.
- The AWS services that are used with the content.
- The country and AWS Region where they store their content.
- The format and structure of their content and whether it is masked, anonymized, or encrypted.
- How their data is encrypted, and where the keys are stored.
- Who has access to their content, and how those access rights are granted, managed, and revoked.

The AWS Shared Responsibility Model also extends to IT controls. The responsibility to operate the IT environment is shared between AWS and its customers, and so is the responsibility for the management, operation, and verification of IT controls. AWS can reduce the administrative load on customers by managing the controls associated with the physical infrastructure deployed in the AWS environment that might previously have been managed by the customer.

AWS Compliance programs

AWS has obtained certifications and independent third-party attestations for a variety of industry-specific workloads. The following compliance programs might be of particular importance to Covered Entities:

- **CSA STAR certification**: The Security Trust Assurance and Risk (STAR) Level 2 certification is a rigorous third-party independent assessment of the security of a cloud service provider. The certification uses the requirements of the ISO/IEC 27001:2013 management system standard together with the CSA Cloud Controls Matrix criteria. The STAR Level 2 certification with STAR validates the use of best practices and the security posture of AWS Cloud offerings. AWS publishes our CSA STAR Level 2 and ISO 27001:2013 certificates on the AWS website and the certificates are also available from AWS Artifact. The Regions and services that are in scope can be found on the CSA STAR Level 2 certification. The AWS services that are in scope for CSA STAR level 2 certification can be found on the ISO and CSA STAR certifications webpage.
- **ISO 27001**: A security management standard that specifies security management best practices and comprehensive security controls that follow the ISO 27002 best practice guidance. The basis of this certification is the development and implementation of a rigorous security program, which includes the development and implementation of an information security management system that defines how AWS perpetually manages security in a holistic, comprehensive manner. For more information or to download the AWS ISO 27001 certification, see the ISO 27001 Compliance webpage.
- **ISO 27017**: Provides guidance on the information security aspects of cloud computing, recommending the implementation of cloud-specific information security controls that supplement the guidance of the ISO 27002 and ISO 27001 standards. This code of practice provides additional implementation guidance for information security controls specific to cloud service providers. For more information or to download the AWS ISO 27017 certification, see the ISO 27017 Compliance webpage.
- **ISO 27018**: Code of practice that focuses on protecting personal data in the cloud. It is based on the ISO information security standard 27002 and provides implementation guidance on ISO 27002 controls that are applicable to cloud personally identifiable information (PII). It also provides a set of additional controls and associated guidance intended to address cloud PII protection requirements not addressed by the existing ISO 27002 control set. For more information or to download the AWS ISO 27018 certification, see the ISO 27018 Compliance webpage.
- **ISO 22301**: Specifies the structure and requirements to implement, maintain, and improve a business continuity management system (BCMS) to protect against, reduce the likelihood of the occurrence of, prepare for, respond to, and recover from disruptions when they arise. Compliance to this standard provides assurance on AWS commitment to business continuity and resiliency of AWS services. For more information or to download the AWS ISO 22301 certification, see the ISO 22301 Compliance webpage.

- **ISO 9001**: Outlines a process-oriented approach to documenting and reviewing the structure, responsibilities, and procedures that are required to achieve effective quality management within an organization. The key to the ongoing certification under this standard is establishing, maintaining, and improving the organizational structure, responsibilities, procedures, processes, and resources so AWS products and services consistently satisfy ISO 9001 quality requirements. For more information or to download the AWS ISO 9001 certification, see the ISO 9001 Compliance webpage.
- **PCI DSS Level 1**: The Payment Card Industry Data Security Standard (PCI DSS) is a proprietary information security standard administered by the PCI Security Standards Council. PCI DSS applies to all entities that store, process, or transmit cardholder data (CHD) or sensitive authentication data (SAD), including merchants, processors, acquirers, issuers, and service providers. The PCI DSS is mandated by the card brands and administered by the Payment Card Industry Security Standards Council. AWS is certified as a PCI DSS Level 1 Service Provider, the highest level of assessment available. For more information or to request the PCI DSS Attestation of Compliance and Responsibility Summary, see the PCI DSS Compliance webpage.
- **SOC**: AWS System and Organization Control (SOC) reports are independent third-party examination reports that demonstrate how AWS achieves key compliance controls and objectives. The purpose of these reports is to help customers and their auditors understand the AWS controls that have been established to support operations and compliance. For more information, see the SOC Compliance webpage. AWS SOC reports come in three forms:
    - **SOC 1**: Provides information about the AWS control environment that might be relevant to a customer's internal controls over financial reporting, in addition to information for the assessment of the effectiveness of internal controls over financial reporting.
    - **SOC 2**: Provides customers and their service users that have a business need with an independent assessment of the AWS control environment relevant to system security, availability, and confidentiality.
    - **SOC 3**: Provides customers and their service users that have a business need with an independent assessment of the AWS control environment relevant to system security, availability, and confidentiality, without disclosing AWS internal information.

See the AWS Compliance Programs webpage for more information about AWS certifications and attestations. See the Best Practices for Security, Identity, and Compliance website for general AWS security controls and service-specific security.

AWS Artifact

Customers can use AWS Artifact to review and download reports and details about more than 2,600 security controls. In addition, AWS Artifact is designed to provide on-demand access to AWS security and compliance documents, including SOC reports, Payment Card Industry (PCI) reports, and certifications from accreditation bodies across geographies and compliance verticals.

AWS Global Cloud Infrastructure

The AWS Global Cloud Infrastructure comprises AWS Regions and Availability Zones. A Region is a physical location in the world that consists of multiple Availability Zones. Availability Zones consist of one or more discrete data centers, each with redundant power, networking, and connectivity, all housed in separate facilities. These Availability Zones offer customers the ability to operate applications and databases, which are more highly available, fault tolerant, and scalable than would be possible in a traditional, on-premises environment. Customers can learn more about these topics by downloading our whitepaper Amazon Web Services Approach to Operational Resilience in the Financial Sector and Beyond.

AWS customers can choose the Region where their content and applications are located. Regions allow AWS customers to establish environments that meet specific geographic or regulatory requirements. Additionally, Regions allow AWS customers with business continuity and disaster recovery objectives to establish primary and backup environments in a location or locations of their choice. More information on our disaster recovery recommendations is available at Disaster Recovery of Workloads on AWS: Recovery in the Cloud.

Further reading

The following resources can help Covered Entities think about security and compliance when designing a secure and resilient environment on AWS.

- AWS Security and Compliance Quick Reference Guide AWS has multiple features to assist in aligning with compliance objectives for regulated workloads on AWS. These features can help achieve a higher level of security at scale. Cloud-based compliance offers a lower cost of entry, simpler operations, and improved agility by providing more oversight, security control, and central automation.
- AWS Security Reference Architecture (AWS SRA) is a holistic set of guidelines for deploying the full complement of AWS security services in a multi-account environment. It can be used to help design, implement, and manage AWS security services so that they align with AWS best practices. The recommendations are built around a single-page architecture that includes AWS security services—how they help achieve security objectives, where they can be best deployed and managed in customer accounts, and how they interact with other security services. This overall architectural guidance complements detailed, service-specific recommendations such as those found on AWS Security Documentation.
- The AWS Well-Architected Framework has been developed to help cloud architects build one of the most secure, high-performing, resilient, and efficient infrastructure possible for their applications. This framework is designed to provide a consistent approach for customers and AWS Partners to evaluate architectures, and provides guidance to help implement designs that scale application needs over time. The AWS Well-Architected Framework consists of six pillars: operational excellence, security, reliability, performance efficiency, cost optimization, and sustainability.
- AWS whitepapers on the six pillars of the AWS Well-Architected Framework: Operational Excellence Pillar, Security Pillar, Reliability Pillar, Performance Efficiency Pillar, Cost Optimization Pillar, and the Sustainability Pillar.
- Global Financial Services Regulatory Principles: AWS has identified five common principles related to financial services regulation that customers should consider when using AWS services and specifically, applying the AWS Shared Responsibility Model to their regulatory requirements. You can review these principles on AWS Artifact.

- NIST Cybersecurity Framework (CSF): The AWS whitepaper NIST Cybersecurity Framework (CSF): Aligning to the NIST CSF in the AWS Cloud shows how organizations can assess the AWS environment against the NIST CSF and improve the security measures they implement and operate (security *in* the cloud). The whitepaper also provides a third-party auditor letter attesting to the conformance to NIST CSF risk management practices (security *of* the cloud) of AWS offerings. Customers can use NIST CSF and AWS resources to support their risk management frameworks.
- AWS Training and Certification: Organizations need individuals with deep security knowledge to help protect their business. AWS training courses enable to build the skillset needed to meet security and compliance objectives. With a variety of content and training materials curated by experts at AWS, customers can stay up to date with evolving best practices and security trends in the industry—whether they are new to the cloud or deployed on AWS.
- Cloud Audit Academy (CAA) is an AWS Security Auditing Learning Path designed for professionals in audit, risk, and compliance roles who assess regulated workloads on AWS. The CAA curriculum includes cloud-specific audit considerations and AWS best practices for security auditing aligned to security and compliance frameworks.

For more information, refer to the Security Learning whitepapers.

Appendix: Considerations on the Second Amendment to NYDFS Cybersecurity Regulation

The final text of the Second Amendment to NYDFS Cybersecurity Regulation is available on this page from the NYDFS. The table that follows aligns each requirement with AWS guidance and AWS services, while mapping the requirements to practices in the AWS Well Architected Framework.

The table is organized into the following columns:

- **Requirements summary**: Summarizes requirements in the Second Amendment to NYDFS Cybersecurity Regulation.
- **Considerations for AWS customers**: Provides considerations for addressing requirements defined in the Second Amendment to NYDFS Cybersecurity Regulation. It outlines AWS services, features, and resources that can help customers address these requirements *in* the cloud. The column also refers to controls implemented and managed by AWS for security and compliance *of* the cloud when applicable.
- **AWS Well-Architected best practices**: Lists best practices for security in the cloud from the AWS Well-Architected Framework that can be implemented as a starting point to support the customer's compliance objectives.

This is not legal or compliance advice. These tables contain only a non-exhaustive sample of considerations and are provided for informational purposes only. Customers are responsible for making their own independent assessments of this information, conducting appropriate due diligence, and should consult with their own legal and compliance advisors.

Requirements summary	Considerations for AWS customers	AWS Well-Architected best practices
500.2 Cybersecurity program **Requirement (c):** Each class A[1] company shall design and conduct independent audits[2] of its cybersecurity program based on its risk assessment. **Compliance date:** April 29, 2024. [1] (d) Class A company means a covered entity with at least $20,000,000 in gross annual revenue in each of the last two fiscal years from all business operations of the covered entity and the business operations in this State of the covered entity's affiliates and: (1) over 2,000 employees averaged over the last two fiscal years, including employees of both the covered entity and all of its affiliates no matter where located; or (2) over $1,000,000,000 in gross annual revenue in each of the last two fiscal years from all business operations of the covered entity and all of its affiliates no matter where located. [2] (h) Independent audit means an audit conducted by internal or external auditors free to make decisions not influenced by the covered entity being audited or by its owners, managers or employees.	**Customer responsibility** This is an action for customers to complete independently. AWS defines some of the most critical aspects of security *in* the cloud for customers through frameworks such as the AWS Well-Architected Framework (which includes a specific Financial Services Industry Lens) and the AWS Cloud Adoption Framework. Both frameworks have specific security areas, including detailed whitepapers, that help focus on how to design and build security features in cloud environments. The following third parties can assist customers with regulatory compliance objectives and assessments: • AWS Security Assurance Services LLC, is a team of industry certified assessors, helping customers achieve, maintain, and automate compliance in the cloud by bringing together applicable audit standards to AWS service specific features and functionality. They help customers build on frameworks such as PCI DSS, HITRUST CSF, NIST, SOC 2, HIPAA, ISO 27001, GDPR, and CCPA. • Security Competency Partners have deep technical expertise with security in AWS and proven customer success securing the cloud journey with their software and services offerings.	Not applicable.

500.3 Cybersecurity policy **Requirements:** Multiple; refer to the text of the Second Amendment to Cybersecurity Regulation for details. **Compliance date:** April 29, 2024	**Customer responsibility** This is an action for customers to complete independently. Changes in this section are related to development, documentation, and implementation of procedures based on the cybersecurity policies and the renaming of cybersecurity policy domain areas. The following third parties can assist customers with regulatory compliance objectives and assessments: • AWS Security Assurance Services LLC, is a team of industry certified assessors, helping customers achieve, maintain, and automate compliance in the cloud by bringing together applicable audit standards to AWS service specific features and functionality. They help customers build on frameworks such as PCI DSS, HITRUST CSF, NIST, SOC 2, HIPAA, ISO 27001, GDPR, and CCPA. • Security Competency Partners have deep technical expertise with security in AWS and proven customer success securing the cloud journey with their software and services offerings.	Not applicable.
500.4 Cybersecurity governance (Chief information security officer) **Requirements:** Multiple; refer to the text of the Second Amendment to Cybersecurity Regulation for details. **Compliance date:** November 1, 2024	**Customer responsibility** This is an action for customers to complete independently. Changes in this section are related to operational requirements for the CISO to report material cybersecurity issues timely, and the roles and responsibility of the senior governing body. Although definition of information security-related roles and responsibilities is an action for customers to complete independently, AWS offers resources and services such as Amazon GuardDuty, AWS Security Hub, and Amazon Security Lake to help customers address these requirements. A common theme among successful AWS customers is that they have an engaged board and senior management team who are both enthusiastic about the benefits of moving to the cloud and are aware of the risks and responsibilities of operating in the cloud. The AWS C-suite Guide to Shared Responsibility for Cloud Security and Data Safe Cloud eBook informs boards and senior management about the benefits and risks of the cloud.	Not applicable.
500.5 Vulnerability management Each covered entity shall, in accordance with its risk assessment, develop and implement written policies and procedures for vulnerability management that are designed to assess and maintain the effectiveness of its cybersecurity program. These policies and procedures shall be	**Customer responsibility** Customers can conduct penetration testing as long as they are limited to the customer's instances and do not violate the AWS Acceptable Use Policy. Advance approval for these scans can be requested through the AWS Penetration Testing Request form. The following third party can assist customers with regulatory compliance objectives and assessments: • Security Competency Partners have deep technical expertise with security in AWS and proven customer success securing the cloud journey with their software and services offerings.	SEC11-BP01 Train for application security SEC11-BP03 Perform regular penetration testing

designed to ensure that covered entities: **Requirement: (a)(1)** conduct, at a minimum, penetration testing of their information systems from both inside and outside the information systems' boundaries by a qualified internal or external party at least annually; **Compliance date**: April 29, 2024		
500.5 Vulnerability management **Requirement: (a)(2)** Conduct, at a minimum automated scans of information systems, and a manual review of systems not covered by such scans, for the purpose of discovering, analyzing and reporting vulnerabilities at a frequency determined by the risk assessment, and promptly after any material system changes. **Compliance date:** May 1, 2025	**Shared responsibility** Under the AWS Shared Responsibility Model, AWS is responsible for vulnerability management of the infrastructure for our managed services. AWS publishes security bulletins to notify customers of security and privacy events with AWS services. AWS also posts information on the list of fixed security vulnerabilities and references to relevant common vulnerabilities and exposures (CVEs) for all supported Amazon Linux versions on Amazon Linux Security Center. Customers are responsible for enabling vulnerability scanning and establishing a process to review and prioritize identified security vulnerabilities on their AWS workloads. Automation is the key to continually scanning workloads for issues and unintended network exposure and performing remediation. AWS has a range of services to help with vulnerability management programs. Amazon Inspector is a comprehensive vulnerability management service that can scan multiple AWS resources, including Amazon Elastic Compute Cloud (Amazon EC2) instances[1], container workloads in Amazon Elastic Container Registry (Amazon ECR)[2], and AWS Lambda functions[3,4]. It identifies different types of vulnerabilities, including software vulnerabilities and unintended network exposure, that can be used to compromise workloads, repurpose resources for unintended use, or exfiltrate data. After it is started, Amazon Inspector automatically discovers all EC2 instances, Lambda functions, and container images residing in Amazon Elastic Container Registry (Amazon ECR) that are identified for scanning, and then scans them for vulnerabilities. Amazon Inspector is designed to provide AWS customers with a detailed list of security findings—prioritized by their severity level—and remediation instructions. Customers can also find a number of vulnerability management solutions in the AWS Marketplace.	SEC06-BP01 Perform vulnerability management SEC06-BP02 Reduce attack surface SEC06-BP03 Reduce manual management and interactive access SEC11-BP02 Automate testing throughout the development and release lifecycle
500.5 Vulnerability management **Requirement: (b)** are promptly informed of new security vulnerabilities by having a monitoring process in place.	**Shared responsibility** AWS Security performs regular vulnerability scans on the host operating system, web application, and databases in the AWS environment using a variety of tools. AWS Security teams also subscribe to newsfeeds for applicable vendor flaws and proactively monitor the vendor websites and other relevant outlets	SEC01-BP04 Keep up-to-date with security threats SEC04-BP02 Capture logs, findings, and

Compliance date: April 29, 2024	for new patches. AWS customers can also report issues to AWS through the AWS Vulnerability Reporting website. Customers are responsible for creating alerts that can notify them of new vulnerabilities on their workloads. Amazon Inspector integration with AWS Security Hub allows AWS customers to send findings from Amazon Inspector to Security Hub. Security Hub collects security data from across AWS accounts, AWS services, and supported third-party partner products, and provides customers with a comprehensive view of their security state in AWS. AWS customers can also use Amazon EventBridge to set up automated alerts to send Amazon Inspector findings to a messaging hub. Amazon Inspector creates an event for EventBridge for newly generated findings, newly aggregated findings, and changes in the state of findings. AWS customers can create an EventBridge rule to send notifications for findings that Amazon Inspector generates to the target specified in the rule. Creating custom responses to Amazon Inspector findings with Amazon EventBridge shows AWS customers how to send alerts for *critical* and *high* severity findings to email, Slack, or Amazon Chime.	metrics in standardized locations SEC04-BP04 Initiate remediation for non-compliant resources SEC10-BP06 Pre-deploy tools
500.5 Vulnerability management **Requirement: (c)** timely remediate vulnerabilities, giving priority to vulnerabilities based on the risk they pose to the covered entity. **Compliance date:** April 29, 2024	**Shared responsibility** AWS Security performs regular vulnerability scans on the host operating systems, web applications, and databases in the AWS environment using a variety of tools. AWS publishes security bulletins to notify customers of security and privacy events affecting AWS services. AWS also posts information on the list of fixed security vulnerabilities and references to relevant CVEs for all supported Amazon Linux versions on Amazon Linux Security Center. AWS Security teams subscribe to newsfeeds for applicable vendor flaws and proactively monitor the vendor websites and other relevant outlets for new patches. AWS customers can also report issues to AWS through the AWS Vulnerability Reporting website. Customers are responsible for patch management for their AWS resources, including EC2 instances, Amazon Machine Images (AMIs), and other compute resources. AWS has a range of services to help with timely remediation of identified vulnerabilities. AWS customers should perform regular patching operations to protect their managed instances from known issues and unauthorized access. AWS customers can use AWS Systems Manager Patch Manager[5] to automate the process of patching EC2 instances managed by AWS Systems Manager using the SSM Agent. Patch Manager, a capability of Systems Manager, automates the process of patching managed nodes with both security-related updates and other types of updates. AWS customers can use Patch Manager to apply patches for both operating systems and applications. A well-designed vulnerability management program should also consider vulnerability testing during the development and deployment stages of the software life cycle. Because implementing vulnerability management during development and deployment helps reduce the probability of a vulnerability making	OPS05-BP05 Perform patch management SEC06-BP04 Automate compute protection REL08-BP01 Use runbooks for standard activities such as deployment REL08-BP04 Deploy using immutable infrastructure REL08-BP05 Deploy changes with automation SUS06-BP02 Keep your workload up-to-date

its way into the production environment. AWS customers can use this guide on automated patching to create an automation solution for patching mutable (long-running) EC2 instances that span multiple AWS accounts and Regions. The guide describes an automated patching solution that uses Lambda to automate patching configurations and scheduling across multiple environments, using Patch Manager and maintenance windows. This solution also uses Amazon QuickSight to provide the necessary reporting and dashboard capabilities to report on patch compliance.

AWS customers might not want to wait for the regular patching schedule to remediate zero-day or other high and critical severity vulnerabilities where patches are available. AWS customers can establish a mechanism to automate on-demand vulnerability management and remediation in their AWS accounts using Amazon Inspector, Security Hub and Systems Manager. The AWS Systems Manager Automation runbook defines the Systems Manager actions on the managed instances and other AWS resources when an automation runs. It contains one or more steps that run in sequential order or branch based on previous steps.

Using Security Hub custom actions, AWS customers can activate a Systems Manager Automation runbook to remediate Amazon Inspector findings for a specific vulnerability identified on multiple EC2 instances[6]. Using this method, AWS customers have the flexibility of selecting Amazon Inspector findings they want to remediate on-demand from Security Hub using custom actions.

AWS customers can also remediate Amazon Inspector findings for EC2 instances at scale within AWS Organizations. AWS customers can patch vulnerabilities identified by Amazon Inspector for their EC2 instances matching a specified tag. AWS customers can also filter Amazon Inspector findings based by severity. This method initiates a Systems Manager Automation runbook from the Systems Manager Automation console[7]. AWS customers can prepare both methods in their AWS account to be ready for use when speed is critical to remediate software vulnerabilities identified by Amazon Inspector.

OS images used to deploy server infrastructure are a key component of reliable and secure compute infrastructure. EC2 Image Builder allows AWS customers to automate the lifecycle of OS images. With Image Builder[8], AWS customers can create images with the necessary components and validate the security of their images before using them in production, reducing their exposure to security vulnerabilities.

AWS customers can configure security scans of their AMIs and container images using Amazon Inspector integration with EC2 Image Builder. After being enabled, Amazon Inspector automatically scans the EC2 instances that Image Builder launches to build and test a new image. Image Builder can generate a security overview of the affected resources, vulnerability details, and known remediations for AMIs and ECS container instances as part of the Image Builder pipeline in their account. See managing security findings for Image Builder images for more information.

500.7 Access privileges and management As part of its cybersecurity program, based on the covered entity's risk assessment each covered entity shall limit user access privileges to information systems that provide access to nonpublic information and shall periodically review such access privileges. **Requirement: (a)(1)** As part of its cybersecurity program, based on the covered entity's risk assessment each covered entity shall limit user access privileges to information systems that provide access to nonpublic information to only those necessary to perform the user's job. **Compliance date:** May 1, 2025	**Customer responsibility** Ensuring that users have appropriate levels of permissions to only access the resources they need is a critical component to enterprise security. The principle of least privilege states that identities should only be permitted to perform the smallest set of actions necessary to fulfill a specific task. This balances usability, efficiency, and security. There are several AWS capabilities to help AWS customers implement the principle of least privilege. **Consideration 1: Authentication and authorization of workforce users to AWS resources** There are multiple ways for human identities to sign in to AWS. The Security Pillar guidance of the Well-Architected Framework states that AWS customers should centralize identity management and aim to reduce reliance on long-term static credentials. It is an AWS best practice to rely on a centralized identity provider using federation when authenticating to AWS. Federation for users can be done either with direct federation to each AWS account or using AWS IAM Identity Center and the identity provider of choice. For federation with individual AWS accounts, AWS customers can use centralized identities for AWS with a SAML 2.0-based provider with AWS Identity and Access Management (IAM). AWS customers can use federation between their AWS account and their chosen provider to grant a user or application access to call AWS API operations by using a SAML assertion to get temporary security credentials. For federation to multiple accounts in their AWS Organizations, AWS customers can configure their identity provider in IAM Identity Center and specify where their users and groups are stored. IAM Identity Center expands the capabilities of IAM to provide a central place that brings together the administration of users and their access to AWS accounts and cloud applications. With IAM Identity Center, AWS customers can define user permissions and manage access to accounts and applications in their AWS Organizations organization centrally. AWS customers can use IAM to implement access control to their AWS resources. AWS customers manage access in AWS by creating policies. A policy is an object in AWS that, when associated with an identity or resource, defines their permissions. AWS supports different types of policies[9]: identity-based policies, resource-based policies, permissions boundaries, AWS Organizations service control policies (SCPs)[10], and session policies. When AWS customers set permissions with IAM policies, they should grant only the permissions required to perform a task by defining the actions that can be taken on specific resources under specific conditions, also known as *least-privilege permissions*. AWS customers might start with broad permissions while they explore the permissions that are required for their workload or use case. As the use case matures, AWS customers can work to reduce the permissions that they grant to work toward least privilege. See grant least privilege access for more information and implementation guidance.	SEC02-BP04 Rely on a centralized identity provider SEC02-BP02 Use temporary credentials SEC03-BP01 Define access requirements SEC03-BP02 Grant least privilege access SEC03-BP05 Define permission guardrails for your organization

	Consideration 2: Manage authentication and fine-grained authorization in applications that AWS customers build User authentication and authorization can be challenging when building web and mobile applications. The challenges include handling authentication, federating identities from external identity providers (IdPs), managing fine-grained permissions, scalability, and more. AWS offers a variety of services to help organizations overcome these challenges. Amazon Cognito is an authentication, authorization, and user management service for web and mobile applications. With Amazon Cognito, AWS customers can authenticate and authorize users from the built-in user directory, from their enterprise directory, and from consumer identity providers such as Google and Facebook. With the launch of Amazon Verified Permissions, AWS customers can use Amazon Cognito as an identity source and add simple, fast authorization to their applications by using Verified Permissions. See Authorization with Amazon Verified Permissions for more information about how to integrate these two services in applications, and an example authorization flow for retrieving resources in applications using Amazon Cognito, Verified Permissions and other supporting AWS services. See Simplify fine-grained authorization with Verified Permissions and Cognito for details. Verified Permissions is a scalable, fine-grained permissions management and authorization service for the applications that AWS customers build. Verified Permissions gives developers and admins a policy and schema management system to define and manage authorization models. Two policy types are supported in Verified Permissions—permit policies and forbid policies—that can be used to explicitly grant or deny access to resources for access and governance control. Policy-based access control helps secure applications without embedding complicated access control code in the application logic. Instead, AWS customers write policies that say who can take what actions on which resources and evaluate the policies by using the Verified Permissions API. AWS also offers a tool for testing and troubleshooting Verified Permissions policies by running a simulated authorization request against policies in the policy store.	
500.7 Access privileges and management **Requirement: (a)** As part of its cybersecurity program, based on the covered entity's risk assessment each covered entity shall: **(2)** Limit the number of privileged accounts and limit the access functions of privileged accounts to only those necessary to perform the user's job;	**Customer responsibility** When an AWS account is created, a root user name and password to sign in to the AWS account are established. The root user of the AWS account has complete access to all AWS services and resources in the account. Root user credentials should not be used for everyday tasks. To view the tasks that require to sign in as the root user, see Tasks that require root user credentials. As a best practice, the root user credentials should be safeguarded as sensitive information. One way to do this is to enable multi-factor authentication (MFA) on the root user. AWS supports three types of MFA devices: virtual MFA devices, FIDO security keys, and hardware MFA devices. For instructions about enabling MFA for root and IAM users in AWS, see enabling	SEC03-BP02 Grant least privilege access SEC03-BP03 Establish emergency access process

(3) Limit the use of privileged accounts to only when performing functions requiring the use of such access. **Compliance date:** May 1, 2025	MFA devices. AWS customers can use MFA devices to safely access multiple AWS accounts, in addition to other token-enabled applications. Read more about the MFA security initiative. AWS recommends creating the necessary IAM roles to perform privileged actions on AWS resources as part of AWS account provisioning process. An IAM role is an object in IAM that is assigned permissions to perform specific actions. AWS customers can use IAM roles to delegate access to users, applications, or services in their AWS account. AWS recommends using identity federation as the preferred way of allowing workforce users[11] into AWS. With identity federation, AWS customers can manage their user identities outside of AWS and grant them access to assume IAM roles based on their job function. AWS customers can configure identity federation on single AWS accounts or can use IAM Identity Center. IAM Identity Center is designed to expand the capabilities of IAM to provide a central place that brings together the administration of user access to AWS accounts and cloud applications. IAM Identity Center reduces the administrative complexity of federating and managing permissions separately for each AWS account. With IAM Identity Center[12], AWS customers can create and manage user identities in the identity store of IAM Identity Center, easily connect to their existing identity source, or another supported IdP. Additionally, AWS IAM permissions policies support conditional access. Using this feature, AWS customers can implement attribute-based access control (ABAC) to define dynamic policies, allowing them to grant permissions based on attributes attached (also known as tags) to IAM principals and AWS resources. This enables them to implement a fine-grained authorization model that combines the advantages of both role-based access controls (RBAC) and ABAC to limit privilege actions to those necessary to perform the job on AWS [13,14]. AWS recommends using automation where possible with no human access to help support audit and compliance needs. However, some cases require human access. For example, unexpected issues might require human intervention to diagnose, fix, or configure manually. For higher-risk scenarios, the organization can supplement baseline access controls by implementing temporary elevated access. The blog post Temporary elevated access management with IAM Identity Center describes a temporary elevated access management solution (TEAM) that integrates with IAM Identity Center and is designed to provide a workflow that allows authorized users to request, review, and approve or reject temporary access. If a request is approved, TEAM activates access for the requester with the scope and duration specified in the request.

500.7 Access privileges and management Requirement: (a)(4) As part of its cybersecurity program, based on the covered entity's risk assessment each covered entity shall: periodically, but at a minimum annually, review all user access privileges and remove or disable accounts and access that are no longer necessary. **Compliance date:** May 1, 2025	**Customer responsibility** Achieving least privilege is a nearly continuous cycle to grant only the permissions that users and systems require. To achieve least privilege, AWS customers can start by setting fine-grained permissions. Then, verify that the existing access meets intent. Finally, refine permissions by removing unused access. The AWS Identity and Access Management Access Analyzer helps AWS customers streamline permissions management. IAM Access Analyzer uses provable security to identify access paths to a resource from outside of its account. It reviews resource policies nearly continuously, and reports findings of public and cross-account access to make it simple to analyze potentially broad access. AWS customers can consider configuring IAM Access Analyzer with AWS Organizations to verify that they have visibility to all accounts. IAM Access Analyzer is designed to let AWS customers preview findings before deploying resource permissions to help validate that their policy changes grant only the intended public and cross-account access to resources. AWS customers can also validate their policies using IAM Access Analyzer policy checks. Create or edit a policy using the AWS Command Line Interface (AWS CLI), AWS API, or JSON policy editor in the IAM console and the IAM Access Analyzer validates the policy against IAM policy grammar and best practices and customers can view policy validation findings that include security warnings, errors, general warnings, and suggestions for their policy. To learn more about validating policies using IAM Access Analyzer, see IAM Access Analyzer policy validation. AWS customers can create dashboards using Amazon QuickSight to visualize IAM Access Analyzer findings and refine permissions. Policy generation is a feature of IAM Access Analyzer that generates a fine-grained policy based on the access activity captured in AWS CloudTrail logs. Using this feature, AWS customers can generate fine-grained policies that grant only the required access. See this overview of how policy generation with IAM Access Analyzer works for a walkthrough of the steps required to generate, customize, and create a fine-grained policy to achieve least privilege.	SEC03-BP04 Reduce permissions continuously SEC03-BP07 Analyze public and cross-account access
500.7 Access privileges and management Requirement: (a)(5) As part of its cybersecurity program, based on the covered entity's risk assessment each covered entity shall: disable or securely configure all protocols that permit remote control of devices **Compliance date:** May 1, 2025	**Customer responsibility** Multiple AWS capabilities can help AWS customers implement the principle of least privilege, such as subnets, security groups, network access control lists (ACLs), AWS Systems Manager Session Manager, and EC2 Instance Connect Endpoint. **Use case: Remote access to AWS compute** Following the principle of least privilege, AWS customers can reduce the attack surface of environments by reducing the set of ports exposed to the network while also restricting the source networks or IP addresses that can have access to EC2 instances. Security groups allow or deny specific inbound and outbound traffic at the resource level such as an EC2 instance. When AWS customers launch an EC2 instance, they can associate it with one or more security groups. In an Amazon Virtual Private Cloud	SEC05-BP01 Create network layers SEC05-BP02 Control traffic at all layers SEC10-BP05 Pre-provision access

(Amazon VPC), EC2 instances run behind a stateful firewall with all ports closed by default. The security group contains rules responsible for opening inbound and outbound ports on that firewall. Although security groups can act as an instance-level firewall, they can also be associated with multiple instances, providing isolation between application tiers in the environment. For example, AWS customers can create a security group for all their web servers that can allow traffic on TCP port 3389, but only from members of the security group containing the remote desktop gateway servers. If customers do not specify a security group when launching an instance, the instance is automatically associated with the default security group for its Amazon VPC.

Network access control lists (ACL) are a set of permissions that AWS customers can attach to a network subnet in a VPC to have stateless filtering of traffic. AWS customers can use network ACLs for inbound or outbound traffic, because they can be an effective way to place a Classless Inter-Domain Routing (CIDR) block or individual IP addresses on a deny list. These ACLs can contain ordered rules to allow or deny traffic based on IP protocol, service port, or source or destination IP address. For example, set a rule that would allow inbound administrative traffic on TCP port 3389 from a specific set of IP addresses. Network ACLs allow or deny specific inbound and outbound traffic at the subnet level.

AWS Systems Manager Session Manager lets AWS customers manage their EC2 instances and on-premises instances through an interactive browser-based shell or the AWS CLI. Session Manager is designed to provide secure instance management without the need to open inbound ports, maintain bastion hosts, or manage Secure Shell (SSH) keys. Session Manager uses the Systems Manager management agent for role-based access to a shell (PowerShell on Windows) on an EC2 instance. Users can access Session Manager from the AWS Management Console or from the command line with the Session Manager AWS CLI plugin. Session Manager is provided at no additional cost for use with EC2 instances. Administrators can use Session Manager to grant and revoke access for both Windows and Linux instances from a single location through IAM roles and policies. There is also no need for privileged account management and password rotation with Session Manager because Session Manager uses IAM roles and policies to check that the user is authorized for the target EC2 instance.

EC2 Instance Connect Endpoint allows AWS customers to connect to an instance through SSH or Remote Desktop Protocol (RDP) without requiring the instance to have a public IPv4 address. After they create an EC2 Instance Connect Endpoint in a subnet, it acts as a private tunnel to the instance. To limit connectivity to only the required instances, AWS recommends configuring a security group for EC2 Instance Connect Endpoint allowing outbound traffic to the specified destination and security groups on the target instances to allow inbound traffic from the endpoint. Access to use EC2 Instance Connect Endpoint is controlled through IAM policies. AWS customers can also define IAM policies to allow users to connect only from a specified source IP address range to control remote access to their instances.

500.7 Access privileges and management Requirement: (b) To the extent passwords are employed as a method of authentication, the covered entity shall implement a written password policy that meets industry standards. **Compliance date:** May 1, 2025	**Customer responsibility** The IAM password policy is an account-level setting that applies to all IAM users, excluding the root user. AWS customers can create a custom password policy on their AWS account to specify complexity requirements such as a minimum password length and specific character types, along with setting mandatory rotation periods for IAM users. These password settings apply only to passwords assigned to IAM users and do not affect the access keys they might have. AWS recommends using identity federation as the preferred way of allowing workforce users into AWS for either console or CLI and API access. With identity federation, AWS customers can manage user identities outside of AWS. It improves security by eliminating individual passwords in each AWS account. AWS customers can leverage the capabilities of external identity sources to implement necessary password complexity requirements.	SEC02-BP05 Audit and rotate credentials periodically SEC02-BP04 Rely on a centralized identity provider
500.7 Access privileges and management Requirement: (c)(1) Each Class A company shall monitor privileged access activity and shall implement: a privileged access management solution. **Compliance date:** May 1, 2025	**Customer responsibility** If a workload requires the storage of credentials necessary to work with other services and resources, AWS Secrets Manager is a service that helps AWS customers securely manage those credentials necessary to their applications, services, and IT resources. A secret can be a password, API key, OAuth token, or other type of privileged credential used for authentication purposes. Using Secrets Manager, AWS customers can secure and manage secrets used to access resources in the AWS environments, on third-party services, and on-premises. This service enables AWS customers to manage, rotate, and retrieve secrets. AWS customers can control access to secrets in Secrets Manager using IAM policies. The blog post Scale your authorization needs for Secrets Manager using ABAC with IAM Identity Center describes a method to define dynamic IAM permissions in IAM Identity Center to control access to secrets based on user attributes from an external identity provider (IdP) and resource tags in Secrets Manager. When using Session Manager to control access to Windows and Linux EC2 instances, there is no need for privileged account management and password rotation. Session Manager uses IAM roles and policies to check that the user is authorized for the target EC2 instance. After being authorized, Session Manager signs the user in with the system-generated OS account, ssm-user. To specify a different OS account, see enable run as support for Linux instances. AWS customers can define dynamic permissions to access EC2 instances with Session Manager in a central place for users federating from an external identity provider using IAM Identity Center and attributes-based access controls. This combination allows AWS customers to control access to specific EC2 instances based on user attributes.	SEC02-BP03 Store and use secrets securely SEC02-BP06 Leverage user groups and attributes

500.7 Access privileges and management **Requirement: (c)(2)** Each Class A company shall monitor privileged access activity and shall implement: an automated method of blocking commonly used passwords for all accounts on the information systems owned or controlled by the class A company and wherever feasible for all other accounts. To the extent the class A company determines that blocking commonly used passwords is infeasible, the covered entity's CISO may instead approve in writing at least annually the infeasibility and the use of reasonably equivalent or more secure compensating controls. **Compliance date:** May 1, 2025	**Customer responsibility** AWS recommends using identity federation as the preferred way of allowing workforce users into AWS for either console or CLI and API access. With identity federation, AWS customers can manage their user identities outside of AWS. Identity federation improves security by eliminating individual passwords in each AWS account. AWS customers can use the capabilities of a supported external identity source to block commonly used passwords.	SEC02-BP01 Use strong sign-in mechanisms SEC02-BP04 Rely on a centralized identity provider
500.8 Application security **Requirement: (b)** All such procedures, guidelines and standards shall be reviewed, assessed and updated as necessary by the CISO (or a qualified designee) of the covered entity at least *annually*. **Compliance date:** April 29, 2024	**Customer responsibility** This is an action for customers to complete independently. The modification brought to this section by the Amendment was to change the frequency of review from periodic to at least annual.	SEC11-BP01 Train for application security SEC11-BP02 Automate testing throughout the development and release lifecycle SEC11-BP03 Perform regular penetration testing SEC11-BP04 Manual code reviews SEC11-BP05 Centralize services for packages and dependencies SEC11-BP06 Deploy software programmatically SEC11-BP07 Regularly assess security properties of the pipelines SEC11-BP08 Build a program that embeds security ownership in workload teams

500.9 Risk assessment **Requirement: (a)** Each covered entity shall conduct a periodic risk assessment of the covered entity's information systems sufficient to inform the design of the cybersecurity program as required by this Part. Such risk assessment shall be reviewed and updated as reasonably necessary, but at a minimum annually, and whenever a change in the business or technology causes a material change to the covered entity's cyber risk. The covered entity's risk assessment shall allow for revision of controls to respond to technological developments and evolving threats and shall consider the particular risks of the covered entity's business operations related to cybersecurity, nonpublic information collected or stored, information systems utilized and the availability and effectiveness of controls to protect nonpublic information and information systems. **Compliance date:** April 29, 2024	**Shared responsibility** AWS performs a nearly continuous risk assessment process to identify, evaluate and mitigate risks across the company. The process involves developing and implementing risk treatment plans to mitigate risks as necessary. AWS monitors and escalates risks, regularly performing risk assessments on newly implemented controls. AWS customers need to update their risk assessment policies, standards, and procedures to reflect the new conditions for reviewing and updating their risk assessments. With regards to material changes in technology, AWS User Notifications can be configured, based on scenarios, to notify the Risk Management organization to review an event to determine if it is a material change to their environment.[15] For more complex scenarios, Amazon EventBridge and Amazon Simple Notification Service (Amazon SNS), can be configured based on specific scoping and filtering criteria. For example: accounts, resources, type of finding, severity of the finding, or tags.[16]	OPS01-BP05 Evaluate threat landscape SEC01-BP03 Identify and validate control objectives SEC01-BP04 Keep up-to-date with security threats SEC01-BP07 Identify threats and prioritize mitigations using a threat model
500.10 Cybersecurity personnel and intelligence **Requirement: (b)** …subject to the requirements set forth in [section] sections 500.4 and 500.11 of this Part. **Compliance date:** April 29, 2024	**Customer responsibility** This is an action for customers to complete independently. The change brought to this section by the Amendment was defining the correlation between section 500.10 – Cybersecurity personnel and intelligence and 500.4 – Chief information security officer.	OPS07-BP01 Ensure personnel capability OPS11-BP04 Perform knowledge management SEC10-BP01 Identify key personnel and external resources SEC11-BP01 Train for application security SEC11-BP08 Build a program that embeds security ownership in workload teams

500.11 Third party service provider security policy **Requirement: (c)** Limited exception. An agent, employee, representative or designee of a covered entity who is itself a covered entity need not develop its own third-party information security policy pursuant to this section if the agent, employee, representative or designee follows the policy of the covered entity that is required to comply with this Part. **Compliance date:** April 29, 2024	**Customer responsibility** The change brought to this section by the Amendment was removing part (c). No additional guidance is necessary.	SEC03-BP09 Share resources securely with a third party
500.12 Multi-factor authentication **Requirement: (a)** Multi-factor authentication shall be utilized for any individual accessing any information systems of a covered entity, unless the covered entity qualifies for a limited exemption pursuant to section 500.19(a) of this Part in which case multi-factor authentication shall be utilized for: **(1)** remote access to the covered entity's information systems; **(2)** remote access to third-party applications, including but not limited to those that are cloud based, from which nonpublic information is accessible **Compliance date:** November 1, 2025	**Customer responsibility** Customers are responsible for configuring multi-factor authentication (MFA) to help protect their AWS resources. **Use case 1: Using MFA for remote access to compute resources** AWS Systems Manager Session Manager lets AWS customers manage their EC2 instances and on-premises instances through an interactive browser-based shell or AWS CLI. Session Manager is designed to provide secure instance management without the need to open inbound ports, maintain bastion hosts, or manage SSH keys. Users can access Session Manager from the AWS Management Console or from the command line with the Session Manager AWS CLI plugin. Securing remote access with MFA outlines a solution that administrators can implement to grant remote access to EC2 instances with multi-factor authentication using Session Manager and IAM Identity Center. Administrators can control access to EC2 instances centrally across multiple accounts using IAM Identity Center. MFA access to EC2 instances can also occur through the existing methods and enterprise directories used in on-premises environments. AWS customers can, of course, implement other systems that enforce MFA access to an operating system such as RADIUS or other third-party directory or MFA token solutions. **Use case 2: Using MFA for cloud applications** With IAM Identity Center, AWS customers can also easily control who can have single sign-on access to cloud applications. Workforce users get one-click access to these applications after they sign in through the IAM Identity Center portal. With IAM Identity Center, AWS customers can enable standard-based strong authentication capabilities for all users across all identity sources. When using IAM Identity Center or Active Directory as the identity source, IAM Identity Center supports the Web Authentication specification to help AWS customers secure user access to AWS	SEC02-BP01 Use strong sign-in mechanisms SEC02-BP02 Use temporary credentials SEC02-BP04 Rely on a centralized identity provider

	accounts and business applications with FIDO-enabled security keys and built-in biometric authenticators, such as facial recognition on computers. AWS customers can also enable one-time-passwords (OTPs) using authenticator apps. For consumer identities, AWS customers can use Amazon Cognito user pools and enable MFA in that service, or by using one of the identity providers that Amazon Cognito user pools supports.	
500.12 Multi-factor authentication **Requirement: (a)(3)** Multi-factor authentication shall be utilized for any individual accessing any information systems of a covered entity, unless the covered entity qualifies for a limited exemption pursuant to section 500.19(a) of this Part in which case multi-factor authentication shall be utilized for: all privileged accounts other than service accounts that prohibit interactive login **Compliance date:** November 1, 2025	**Customer responsibility** Customers are responsible for configuring multi-factor authentication (MFA) to help protect their AWS resources. AWS Multi-factor Authentication (MFA) for IAM for highly privileged users is a security feature that augments user name and password credentials. MFA requires users to prove physical possession of a hardware MFA token or MFA-enabled mobile device by providing a valid MFA code. Customers can enable MFA at the AWS account level and for the root user and IAM users they have created in their account. IAM permissions policies support conditional access. Using this feature, AWS customers have the ability to create IAM policies with conditional access based on MFA usage, enabling them to implement additional layers of authentication for privileged actions on AWS. For users federating from an external identity provider, conditional access based on MFA can be implemented using SAML session tags provided by the identity provider, and passed through the SAML assertion to be consumed and validated by AWS. Session tags are key-value pair attributes passed when a user assumes an IAM role. AWS customers can pass SAML attributes as session tags with the AssumeRoleWithSAML operation, and use PrincipalTag key in the condition element of policies to allow or deny access based on MFA authentication response from IdP. For the AWS account root user, AWS supports three types of MFA devices: virtual MFA devices, FIDO security keys, and hardware MFA devices. Virtual MFA devices are software-based apps, usually running on a mobile device, that generate secure, one-time authentication codes that are used as part of the sign-on process. FIDO security keys and hardware MFA devices are physical devices that are required to gain access to the accounts to which they are attached. These physical devices are considered some of the most secure options for MFA. AWS customers can register up to eight MFA devices of any combination of the currently supported MFA types with their AWS account root user and IAM users. For instructions about enabling MFA for the root user and IAM users in AWS, see enabling MFA devices. Because MFA is so important, AWS can provide an MFA device at no cost to qualified account holders. AWS customers can use MFA devices to safely access multiple AWS accounts, in addition to other token-enabled applications. Read more about the MFA security initiative.	SEC01-BP02 Secure account root user and properties SEC03-BP01 Define access requirements SEC03-BP02 Grant least privilege access

500.13 Asset management and limitations on data retention requirements **Requirement: (a)** "…implement written policies and procedures designed to produce and maintain a complete, accurate and documented asset inventory of the covered entity's information systems. The asset inventory shall be maintained in accordance with written policies and procedures. At a minimum, such policies and procedures shall include: a method to track key information for each asset, including, as applicable, the following: ▪ owner; ▪ location; ▪ classification or sensitivity; ▪ support expiration date; and ▪ recovery time requirements" **Compliance date:** November 1, 2025	**Customer responsibility** AWS customers can use the following AWS services and resources to assist them[17]: AWS Config is designed to work with a detailed inventory of the customer's resources on AWS and their configuration and to record resource configuration changes, including how resources are related to one another and how they were configured in the past, so that AWS customers can see the configurations and relationships change over time[18]. Amazon CloudWatch is designed to provide data and actionable insights to monitor applications, understand and respond to system-wide performance changes, improve resource utilization, and get a unified view of operational health. AWS Systems Manager can assist with visibility and control of customer infrastructure on AWS. In addition, AWS Systems Manager is designed to view operational data from multiple AWS services and automate operational tasks across AWS resources. AWS Service Catalog is designed to allow creating and managing catalogues of IT services and infrastructure resources approved for use on AWS. In addition, AWS customers can define and manage relationships between these resources, enabling proper change and configuration management processes. AWS resource tagging is how AWS customers can manage, identify, organize, search, and filter resources by assigning metadata to their AWS resources. For example, security tags can contain information on confidentiality, identifying the specific data confidentiality levels that a resource supports, or compliance, such as an identifier for workloads that must adhere to specific compliance requirements. AWS customers can use AWS Resource Groups to organize their AWS resources. AWS Resource Groups are designed to help manage and automate tasks on large numbers of resources. In addition, customers can assign metadata to resources using tags.	OPS04-BP02 Implement application telemetry OPS04-BP03 Implement user experience telemetry OPS05-BP03 Use configuration management systems OPS08-BP03 Analyze workload traces OPS08-BP04 Create actionable alerts OPS09-BP03 Review operations metrics and prioritize improvement OPS10-BP06 Communicate status through dashboards OPS10-BP07 Automate responses to events OPS11-BP07 Perform operations metrics reviews SEC 1 Secure Operations SEC 2 Authentication SEC 3 Authorization and access control SEC 4 Security events SEC06-BP01 Perform vulnerability management SEC06-BP04 Automate compute protection
500.14 Monitoring and training As part of its cybersecurity program, each covered entity shall: **Requirement: (a)(2)** implement risk-based controls designed to protect against malicious code, including those that monitor and filter web traffic and electronic mail to block malicious content. **Compliance date:** May 1, 2025	**Customer responsibility** AWS customers can use the following AWS services and resources to assist them: **Use case 1: External traffic monitoring** AWS Web Application Firewall (AWS WAF) is a web application firewall that helps protect web applications from attacks by giving AWS customers functionality to configure rules that allow, block, or monitor (count) web requests based on conditions that they define. These conditions include IP addresses, HTTP headers, HTTP body, URI strings, SQL injection and cross-site scripting. AWS also works directly with third parties providing AWS WAF products that are vetted by AWS Partner Solutions architects. For	SEC01-BP07 Identify threats and prioritize mitigations using a threat model SEC04-BP03 Correlate and enrich security alerts SEC05-BP01 Create network layers SEC05-BP02 Control traffic at all layers

	more information, see Protect Your Web Applications with AWS WAF Ready Partners. Amazon Simple Email Service (Amazon SES) lets AWS customers reach their customers without an on-premises Simple Mail Transfer Protocol (SMTP) system. Amazon SES uses multiple spam and virus protection measures. It uses block lists to prevent mail from known spammers from entering the system. It also performs virus scans on incoming email that contains an attachment. Amazon SES makes its spam detection verdicts available to AWS customers, enabling them to decide if they trust each message. In addition to the spam and virus verdicts, Amazon SES is designed to provide the DomainKeys Identified Mail (DKIM) and sender policy framework (SPF) check results. Organizations using third party email products can find email security monitoring and gateway solutions within AWS Marketplace. **Use case 2: Internal network, compute, storage, and database monitoring** Amazon GuardDuty[19] is a threat detection service that nearly continuously monitors AWS accounts, EC2 instances, Lambda functions, Amazon Elastic Kubernetes Service (Amazon EKS) clusters, Amazon Aurora sign-in activity, and data stored in Amazon S3 for unexpected activity. If GuardDuty detects anomalous behaviors such as credential exfiltration, or command and control infrastructure (C2) communication, GuardDuty will generate detailed security findings that can be used for security visibility and remediation. Additionally, Amazon GuardDuty EKS Runtime Monitoring is designed for runtime threat detection coverage for Amazon EKS nodes and containers within the AWS environment. Findings from the GuardDuty EKS Runtime Monitoring provide metadata context to identify potential threats and contain them before they escalate. GuardDuty EKS Runtime Monitoring uses a fully managed GuardDuty security agent that adds visibility into individual container runtime activities, such as file access, processes, and network connections. The GuardDuty security agent helps GuardDuty identify specific containers within EKS clusters that are potentially compromised. It can also detect attempts to escalate privileges from an individual container to the underlying EC2 instance, and the broader AWS environment.	SEC05-BP03 Implement inspection-based protection SEC05-BP04 Automate network protection SEC10-BP06 Pre-deploy tools PERF05-BP01 Understand how networking impacts performance
500.14 Monitoring and training **Requirement: (b)(1)** Each class A company shall implement, unless the CISO has approved in writing the use of reasonably equivalent or more secure compensating controls: an endpoint detection and response solution to monitor anomalous activity, including	**Customer responsibility** The introduction of distributed IT environments, mobile devices, or IoT devices has expanded an organization's attack surface. Monitoring across these endpoints[20] for anomalous activity performed by users is an important part of a cybersecurity program. AWS customers can use the following AWS services and resources to assist them: **Use case 1: Comprehensive endpoint scanning** Endpoint detection and response (EDR) software has advanced risk detection, investigation, and remediation capabilities. It is an endpoint security solution that almost continuously monitors end-	SEC06-BP01 Perform vulnerability management SEC06-BP02 Reduce attack surface SEC06-BP03 Reduce manual management and interactive access SEC11-BP02 Automate testing throughout the development and release lifecycle

but not limited to lateral movement. **Compliance date:** November 1, 2024	user devices to more quickly detect and respond to security incidents. EDR works by doing the following: • Records activities and events taking place on all endpoints. • Analyzes events in real time to detect suspicious behavior automatically. • Provides continuous and comprehensive visibility into what is happening on endpoints in real time. EDR tools provide intelligence that security teams can use to proactively investigate, minimize, and respond to security risks. Endpoint solutions available in the AWS Marketplace help AWS customers manage and configure endpoint assets and secure them against bugs, malware, and inadvertent data disclosure. **Use case 2: Amazon GuardDuty malware detection** Amazon GuardDuty Malware Protection[21] helps AWS customers detect the potential presence of malware by scanning the Amazon Elastic Block Store (Amazon EBS) volumes that are attached to EC2 instances and container workloads. Malware Protection offers two types of scans to detect unexpected activity in EC2 instances and container workloads: a malware scan initiated by GuardDuty, and an on-demand malware scan. After customers enable malware scans initiated by GuardDuty, GuardDuty will generate findings if it detects activity that indicates the potential presence of malware in an EC2 instance or container workload, and will automatically initiate an agentless scan on the EBS volumes attached to the potentially impacted EC2 instance or container workload to detect the presence of malware. AWS customers can initiate an on-demand malware scan by providing the Amazon Resource Name (ARN) of the EC2 instance that they want to scan through the GuardDuty console or API. As a response to an on-demand malware scan or an automatically invoked malware scan initiated by GuardDuty, GuardDuty creates snapshots of the relevant EBS volumes attached to the potentially impacted resources, and shares them with the GuardDuty service account. Malware Protection creates encrypted EBS volumes from those snapshots in the service account. After the scan is complete, GuardDuty deletes the encrypted EBS volumes and the snapshots of EBS volumes. If malware is found and the snapshot retention setting is on, the snapshots of the EBS volumes cannot be deleted and are automatically retained in the AWS account. If no malware is found, the snapshots of EBS volumes are not be retained, regardless of the snapshot retention setting. By default, the snapshots retention setting is off.	
500.14 Monitoring and training **Requirement: (b)(2)** Each class A company shall implement, unless the CISO has approved in writing the use of reasonably equivalent or more secure compensating controls: a	**Customer responsibility** An organization's cybersecurity and operational risks increase when the Security Operation Center (SOC) is challenged to identify cybersecurity issues timely, because of numerous and disparate logging sources, growing log volumes, and unique log formats or schemas. **Use case 1: Centralized logging**	SEC04-BP01 Configure service and application logging SEC04-BP02 Capture logs, findings, and metrics in standardized locations

solution that centralizes logging and security event alerting.
Compliance date: November 1, 2024

Amazon Security Lake is a service designed to automatically centralize security logs and findings from AWS environments, software as a service (SaaS) providers, on premises, and cloud sources into a purpose-built data lake stored in the customer account allowing AWS customers to retain control and ownership over their security logs.

Security Lake can help make the organization's security data broadly accessible to preferred security analytics solutions to assist in use cases such as threat detection, investigation, and incident response. Security engineers can get broad visibility to investigate and respond to security events.

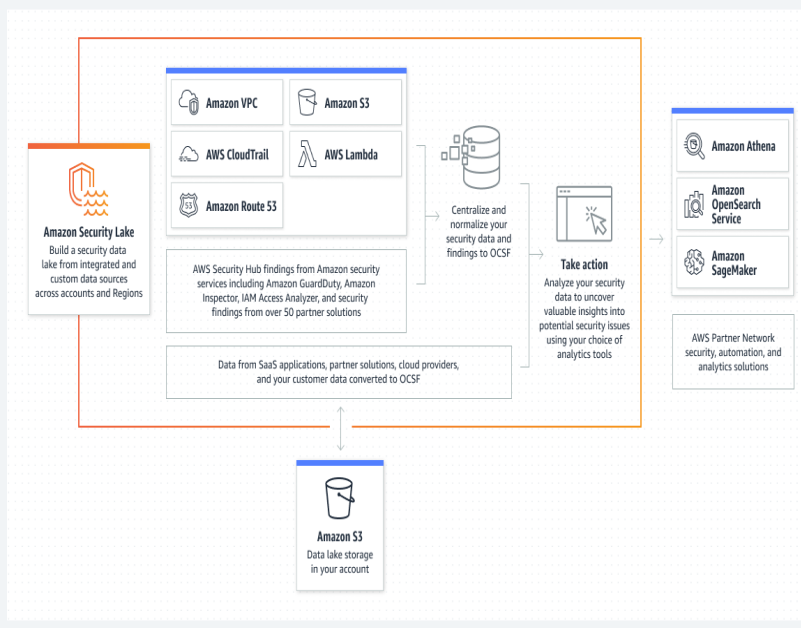

Figure 2 – Security Lake architecture.

Security Lake has adopted the Open Cybersecurity Schema Framework (OCSF), an open standard. With OCSF support, the service normalizes and combines security data from AWS and a broad range of enterprise security data sources.

For additional information about Security Lake, see Amazon Security Lake is now generally available | AWS Security Blog.

Use case 2: Security event alerts with Security Lake

Security Lake offers integration with third-party security services such as security information and event management (SIEM) and extended detection and response (XDR) tools to support security event alerting.

Use case 3: Security data analytics

Beyond centralized security data management and security event alerting, Security Lake also offers integration with AWS data analytics and machine learning services. Amazon OpenSearch Service, Amazon OpenSearch Ingestion, and Amazon SageMaker can be configured to analyze and visualize petabytes of data to uncover potential security issues. For more information, see:

SEC04-BP03 Correlate and enrich security alerts
SEC04-BP04 Initiate remediation for non-compliant resources
SEC10-BP02 Develop incident management plans
SEC10-BP03 Prepare forensic capabilities
OPS07-BP04 Use playbooks to investigate issues

	• Generate security insights from Amazon Security Lake data using Amazon OpenSearch Ingestion • Ingest, transform, and deliver events published by Amazon Security Lake to Amazon OpenSearch Service • Generate security insights from Amazon Security Lake data using Amazon OpenSearch Ingestion • Generate machine learning insights for Amazon Security Lake data using Amazon SageMaker	
500.15 Encryption of nonpublic information **Requirement: (a)** As part of its cybersecurity program, each covered entity shall implement a written policy requiring encryption that meets industry standards, to protect nonpublic information held or transmitted by the covered entity both in transit over external networks and at rest. **Compliance date:** November 1, 2024	**Customer responsibility** AWS provides ways to categorize data based on levels of sensitivity, such as nonpublic information. By using resource tags, IAM policies, and Amazon Macie, customers can define and implement policies for data classification. **Use case 1: Encryption at rest** AWS is committed to provide industry standard encryption services and capabilities to protect both the cloud infrastructure used by our customers and for our customers to use within their environments. Third party auditors have assessed AWS compliance with industry standards including NIST FIPS140-2, SOC, PCI, FedRAMP, HIPAA, and others. For more information regarding third party assessments see AWS Artifact. AWS offers encryption for services where data is stored[22] and the ability to add layers of security to data at rest in the cloud, providing scalable and efficient encryption features.[23] AWS Key Management Service (AWS KMS) is designed to help AWS customers create and manage cryptographic keys (for example, 256-bit AES-GCM[24,25]) and control their use across multiple AWS services and in the customer's applications. AWS KMS is a secure and resilient service that uses hardware security modules that have been validated under FIPS 140-2 Level 3, or are in the process of being validated, to protect keys. AWS asserts as a fundamental security principle that there is no human interaction with plaintext cryptographic key material in AWS services. There is no mechanism for anyone, including AWS service operators, to view, access, or export plaintext key material. AWS KMS is also integrated with AWS CloudTrail to give AWS customers logs to help meet their regulatory and compliance needs. For the complete list of AWS services integrated with AWS KMS, see AWS Service Integration. AWS CloudHSM is a cloud-based hardware security module (HSM) that is designed to enable AWS customers to generate and use their own encryption keys on AWS. With CloudHSM, AWS customers can manage their own encryption keys using single-tenant FIPS 140-2 Level 3 validated HSMs. CloudHSM offers AWS customers the flexibility to integrate with their applications using industry-standard APIs, such as PKCS#11, Java Cryptography Extensions (JCE), and Microsoft CryptoNG (CNG) libraries. The AWS Encryption SDK is designed to provide a client-side encryption library for implementing encryption and decryption operations on all types of data.	SEC02-BP03 Store and use secrets securely SEC08-BP01 Implement secure key management SEC08-BP02 Enforce encryption at rest SEC08-BP03 Automate data at rest protection SEC09-BP02 Enforce encryption in transit

AWS Secrets Manager is designed to help AWS customers protect secrets needed to access their applications, services, and IT resources. The service enables AWS customers to rotate, manage, and retrieve database credentials, API keys, and other secrets. Users and applications retrieve secrets with a call to Secrets Manager APIs, eliminating the need to hardcode sensitive information in plain text. AWS Secrets Manager offers secret rotation with built-in integration for Amazon Relational Database Service (Amazon RDS), Amazon Redshift, and Amazon DocumentDB. Also, the service is extensible to other types of secrets, including API keys and OAuth tokens. All secrets stored in Secrets Manager are encrypted by default using keys backed by AWS KMS.

Use case 2: Encryption in transit

To protect data in transit, AWS encourages customers to use a multi-level approach. All network traffic between AWS data centers is transparently encrypted at the physical layer. All traffic within a VPC and between peer VPCs across Regions is transparently encrypted at the network layer when using supported EC2 instance types. At the application layer, customers have a choice about whether and how to use encryption using a protocol such as Transport Layer Security (TLS). All AWS service endpoints support TLS to create a secure HTTPS connection to make API requests.

AWS Certificate Manager (ACM) is a service that AWS customers can use to provision, manage, and deploy public and private Secure Sockets Layer/Transport Layer Security (SSL/TLS) certificates for use with AWS services and internal connected resources.

AWS Private Certificate Authority (AWS Private CA) is designed to provide AWS customers with a highly-available private Certificate Authority (CA) without the upfront investment and ongoing maintenance costs of operating their own private CA. [26]

The AWS Encryption SDK is designed to provide a client-side encryption library for implementing encryption and decryption operations on all types of data.

AWS Direct Connect is designed to be an alternative to using the internet to connect to AWS. Using Direct Connect, data is delivered through a private network connection between customer facilities and AWS. In many circumstances, private network connections can reduce costs, increase bandwidth, and provide a more consistent network experience than internet-based connections. All AWS services, including Amazon EC2, Amazon VPC, Amazon S3, and Amazon DynamoDB can be used with Direct Connect.

AWS VPN is designed to offer secure and private access to resources with either an AWS Site-to-Site VPN, Accelerated Site-to-Site VPN, or Client VPN connection.

AWS PrivateLink enables customers to access services hosted on AWS in a highly available and scalable manner, while keeping all the network traffic within the AWS network. Service users can privately access services powered by AWS PrivateLink from their Amazon VPC or their on-premises data centers, without using

	public IPs, and without requiring traffic to traverse across the internet.	
500.15 Encryption of nonpublic information **Requirement: (b)** To the extent a covered entity determines that encryption of nonpublic information at rest is infeasible, the covered entity may instead secure such nonpublic information using effective alternative compensating controls that have been reviewed and approved by the covered entity's CISO in writing. The feasibility of encryption and effectiveness of the compensating controls shall be reviewed by the CISO at least annually. **Compliance date:** November 1, 2024	**Customer responsibility** Encryption maintains the confidentiality of sensitive data in the event of unauthorized access or accidental disclosure. AWS recommends enforcing the use of encryption for data at rest. AWS KMS integrates with AWS services to encrypt data at rest. For example, AWS customers can set default encryption on an Amazon S3 bucket so that new objects are automatically encrypted. Similarly, AWS customers can enable default encryption at the account level so that Amazon EBS encrypts the EBS volumes created when they launch a new EC2 instance, in addition to new snapshot copies that they create. However, to help protect data at rest, AWS also recommends implementing multi-layered controls including, but not limited to, enforcing access using mechanisms such as isolation and versioning, and applying the principle of least privilege. • Separate data based on different classification levels and verify that only authorized users can access data on a need-to-know basis. Implement mechanisms such as automation to keep people away from data. • Protect data with regular backups and versioning to prevent intentional or inadvertent modification or deletion of data. Use Amazon S3 versioning and object lock when appropriate. • Access to data should be audited using detective mechanisms, such as CloudTrail, and service level logs, such as Amazon S3 access logs.	SEC07-BP03 Automate identification and classification SEC04-BP02 Capture logs, findings, and metrics in standardized locations REL09-BP02 Secure and encrypt backups
500.16 Incident response and business continuity management **Requirement: (a)(1)(vii)** As part of its cybersecurity program, each covered entity shall establish written plans that contain proactive measures to investigate and mitigate cybersecurity events and to ensure operational resilience, including but not limited to incident response, business continuity and disaster recovery plans. Incident response plans shall be reasonably	**Customer responsibility** Establishing organizational plans is an action for AWS customers to complete independently. AWS customers can engage the AWS customer incident response team (CIRT) through a support case to assist with triage and recovery for an active security event on AWS. The AWS CIRT is a specialized global AWS team that supports customers during active security events on the customer side of the AWS Shared Responsibility Model. The team is made up of AWS Global Services Consultants and Solutions Architects with experience in incident response. When engaged, AWS CIRT assists in root-cause analysis through AWS service logs and can give recommendations for recovery. In addition, AWS CIRT can give security tips and best practices to help customers avoid security events in the future. Backup and restore processes help restore data to a point in time before data corruption, modification, or destruction. These data events can be accidental or part of a cybersecurity event.	OPS10-BP01 Use a process for event, incident, and problem management OPS11-BP02 Perform post-incident analysis SEC10-BP07 Run simulations REL09-BP01 Identify and back up all data that needs to be backed up, or reproduce the data from sources REL09-BP02 Secure and encrypt backups

designed to enable prompt response to, and recovery from, any cybersecurity event… …including disruptive events such as ransomware incidents… …recovery from backups… **Compliance date:** November 1, 2024	Implementing backup and restore processes can help AWS customers prepare to respond and recover from these events with acceptable recovery time objectives (RTOs) and recovery point objectives (RPOs). AWS provides several solutions for backups to integrate with operational and security incident recovery procedures. For more information on backup capabilities for Amazon EC2, transactional databases, Amazon Elastic File System (Amazon EFS), and Amazon S3, see Use backups to recover from security incidents. **Use case 1: Backup from on-premises to AWS Cloud** AWS Storage Gateway is a cloud storage solution that helps customers overcome cloud storage challenges and bridge the gap between their on-premises environments and the cloud. AWS Storage Gateway enables on-premises applications to use cloud storage by providing low-latency data access over standard storage protocols.[27] AWS Backup can be used to back up on-premises AWS Storage Gateway volumes and VMware virtual machines, providing a common method to manage the backups of application data both on premises and on AWS. AWS Backup defines a central data protection policy called a backup plan that works across AWS services for compute, storage, and databases. The backup plan defines parameters such as backup frequency and backup retention period. After AWS customers define their data protection policies and assign AWS resources to the policies, AWS Backup automates the creation of backups and stores those backups in an encrypted backup vault that AWS customers designate.[28] **Use case 2a: Backup within AWS Cloud (Amazon S3 only)** Based on their NYDFS risk assessment, AWS customers might determine that specific workloads only require data storage to be backed up. Amazon S3 is an object storage service offering scalability, data availability, security, and performance and using Amazon S3 Versioning, AWS customers can keep multiple versions of an object in the same Amazon S3 bucket, which gives them the ability to restore a particular version during the recovery process. AWS customers can use Amazon S3 Versioning to preserve, retrieve, and restore versions of objects stored in their Amazon S3 buckets. With versioning, they can recover from both unintended user actions and application failures, such as accidental deletion or overwrite. [29] **Use case 2b: Backup within AWS (compute, storage, and database)** AWS Backup[30] defines a central data protection policy called a backup plan that works across AWS services for compute, storage, and databases. The backup plan defines parameters such as backup frequency and backup retention period. After AWS customers define their data protection policies and assign AWS resources to the policies, AWS Backup automates the creation of backups and stores those backups in an encrypted backup vault that they designate.[31] For more information, see AWS Backup Blogs. AWS Backup Audit Manager can almost continuously evaluate backup activity and generate audit reports that can help AWS	REL09-BP03 Perform data backup automatically REL09-BP04 Perform periodic recovery of the data to verify backup integrity and processes

	customers demonstrate compliance with regulatory requirements. These reports also provide AWS customers with more visibility into backup activities, helping them monitor their operational posture and identify failures that might need further action. AWS Elastic Disaster Recovery (AWS DRS) can be used for ransomware recovery. Elastic Disaster Recovery can launch unlocked and unencrypted versions of servers from before the ransomware event into a preferred Region. This point-in-time recovery capability protects data and enables AWS customers to be up and running minutes after a ransomware event—without having to pay ransom.	
500.16 Incident response and business continuity management **Requirement: (a)(1)(viii)** …establish written plans that contain proactive measures to investigate and mitigate cybersecurity events and to ensure operational resilience... …preparation of root cause analysis that describes how and why the event occurred, what business impact it had, and what will be done to prevent reoccurrence. **Compliance date:** November 1, 2024	**Customer responsibility** AWS customers can use Amazon Detective to analyze, investigate, and identify the root cause of security findings or suspicious activities. Amazon Detective automatically collects log data from their AWS resources. It uses machine learning, statistical analysis, and graph theory to generate visualizations that help AWS customers conduct security investigations.	SEC04-BP01 Configure service and application logging SEC04-BP02 Capture logs, findings, and metrics in standardized locations SEC10-BP03 Prepare forensic capabilities OPS07-BP04 Use playbooks to investigate issues
500.16 Incident response and business continuity management **Requirement: (a)(2)** Business continuity and disaster recovery (BCDR) plan. BCDR plans shall be reasonably designed to ensure the availability and functionality of the covered entity's information systems and material services and protect the covered entity's personnel, assets and nonpublic information in the event of a cybersecurity-related disruption to its normal business activities…" **Compliance date:** November 1, 2024	**Customer responsibility** Disaster recovery (DR) is a necessary aspect of resiliency strategies in the cloud as it defines how workloads respond when a disaster—defined as an event that causes a significant negative impact on the business—occurs. To address DR, organizations must consider their business objectives and implement resilience in the design of their workloads to meet their recovery objectives, known as the recovery point objective (RPO) and recovery time objective (RTO), for a one-time disaster event. AWS offers disaster recovery services that can help customers replicate their data and applications to a secondary location, enabling them to recover in the event of a disaster. These services include AWS Storage Gateway, Amazon S3, Amazon EBS, and AWS Elastic File System (Amazon EFS). AWS also provides AWS Elastic Disaster Recovery, a fully managed service that customers can use to set up disaster recovery solutions. AWS Elastic Disaster Recovery automates the replication of data and applications and is designed to provide continuous monitoring and testing of the disaster recovery environment. Customers might consider deploying infrastructure across multiple AWS Regions for highly critical workloads with data replication and continuous backups to minimize business impact. Disaster recovery on AWS offers several advantages over traditional	REL09-BP01 Identify and back up all data that needs to be backed up, or reproduce the data from sources REL09-BP03 Perform data backup automatically REL09-BP04 Perform periodic recovery of the data to verify backup integrity and processes REL10-BP02 Select the appropriate locations for your multi-location deployment REL10-BP03 Automate recovery for components constrained to a single location

	environments, such as reduced complexity, repeatable testing, lower management overhead, automation opportunities, and reduced cost. **Use case 1: Disaster recovery from on-premises to AWS** For information about using AWS as a disaster recovery site for on-premises workloads, see: Disaster Recovery of On-Premises Applications to AWS. **Use case 2: Disaster recovery within AWS**[32] An event that prevents a workload or system from fulfilling its business objectives in its primary deployed location is considered a disaster. Disaster Recovery of Workloads on AWS: Recovery in the Cloud outlines the best practices for planning and testing disaster recovery for workloads deployed to AWS, and offers different approaches to mitigate risks and meet the Recovery Time Objective (RTO) and Recovery Point Objective (RPO) for that workload.	REL13-BP01 Define recovery objectives for downtime and data loss REL13-BP02 Use defined recovery strategies to meet the recovery objectives REL13-BP03 Test disaster recovery implementation to validate the implementation REL13-BP04 Manage configuration drift at the DR site or Region REL13-BP05 Automate recovery
500.16 Incident response and business continuity management **Requirement: (d)(2)** Each covered entity shall periodically, but at a minimum annually, test its: …Ability to restore its critical data and information systems from backups. **Compliance date:** November 1, 2024	**Customer responsibility** The following AWS services can help customers comply with this requirement by assisting in testing the recovery of data. **Use case 1: Failover testing from on-premises to AWS** The AWS on-demand and pay as you go pricing allows AWS customers to perform restoration testing from on-premises to the cloud and only use the resources they need. For more information, read Testing AWS Direct Connect Resiliency with Resiliency Toolkit – Failover Testing. **Use case 2: Data recovery within AWS** A backup strategy must include testing backup plans. A backup strategy is not effective if the data in the backups cannot be restored. To increase confidence in the ability to recover backup data, AWS customers can create a basic and repeatable process for continuous data recovery testing.[33] AWS customers can automate recovery testing by creating a pipeline to validate backups using AWS Backup, Amazon EventBridge, and AWS Lambda. Having an automated solution to validate data recovery from backups can increase confidence in using these data backups for disaster recovery. For more information, read the AWS blog: Automate data recovery validation with AWS Backup. **Use case 3: IT system recovery within AWS (at scale)** AWS Elastic Disaster Recovery can be used to perform non-disruptive DR testing. AWS Elastic Disaster Recovery enables a mechanism for AWS customers to test recovery of their source environment at scale without performance impact. Testing the DR plan regularly helps AWS customers verify that they can meet recovery objectives and recover applications after an IT disruption.[34] **Use case 4: Advanced recovery testing within AWS** It is necessary to assess and test the disaster recovery plan regularly, and AWS Resilience Hub can help by monitoring the resilience of workloads.	REL09-BP04 Perform periodic recovery of the data to verify backup integrity and processes REL13-BP03 Test disaster recovery implementation to validate the implementation REL12-BP05 Test resiliency using chaos engineering

	Workloads architected with automated failover and recovery processes can use AWS Fault Injection Service to run tests to confirm processes operate as designed. The AWS Fault Injection Service is a fully managed service for running fault injection experiments to improve the performance, observability, and resiliency of an application. It supports the process of setting up and running controlled fault injection tests across a range of AWS services, so teams can build confidence in their application behavior.	
500.16 Incident response and business continuity management **Requirement: (e)** Each covered entity shall maintain backups necessary to restore material operations. The backups shall be adequately protected from unauthorized alterations or destruction. **Compliance date:** November 1, 2024	**Customer responsibility** AWS can provide data protection and data vault solutions to help customers address the risks of unauthorized alteration or destruction of backups. These solutions can be deployed at scale, with faster time to value, and at a reduced cost compared to on-premises solutions. AWS can work with third parties that specialize in data protection to support customers in extending their workflows and operational procedures into the cloud to drive performance, efficiency and scale. **Use case 1: Storage level protection** S3 Object Lock is used for object storage immutability and ransomware protection within cloud storage, backup, and data protection solutions. In *compliance mode*, a protected object version cannot be overwritten or deleted by any user, including the AWS account root user. When an object is locked in compliance mode, the retention mode cannot be changed, the retention period cannot be shortened, and an object version cannot be overwritten or deleted for the duration of the retention period. *Using the Cyber Vault Environment Patterns* AWS customers can protect their backups from unauthorized alterations or destruction by building a cyber vault environment in alignment with their NYDFS risk assessment using the Cyber Vault Environment Patterns. For lower risk data, a cyber vault can be implemented using secure standard configurations. For higher risk data, a cyber vault can be customized using a combination of network isolation or air-gapping[35], access controls, forensic analysis, and malware detection. 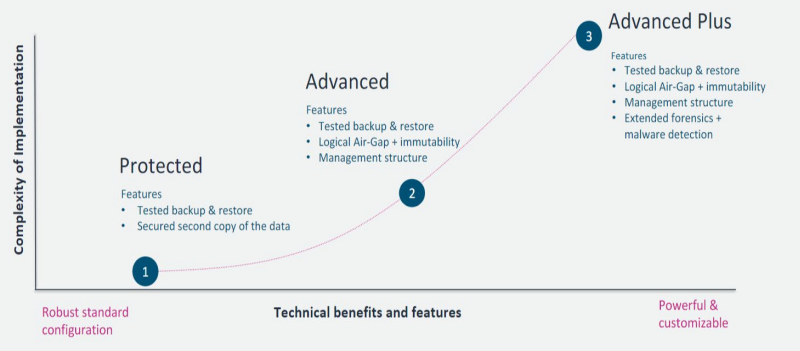 *Figure 3 – Cyber vault environment patterns.*	REL09-BP02 Secure and encrypt backups SEC08-BP02 Enforce encryption at rest SEC08-BP03 Automate data at rest protection SEC08-BP04 Enforce access control SEC08-BP05 Use mechanisms to keep people away from data

Phase 1: The *Protected* vault pattern is a secure data vault environment used to protect critical data backups of an organization in a separate high security area. The data is stored using immutable WORM (write once, read many) storage.

Phase 2: The *Advanced* vault pattern is logically air-gapped to prevent unintended access to the data vault. The vault is configured to periodically copy data from the source into WORM storage. In this case, the data vault is kept isolated and mostly unreachable from the outside using firewalls and zero trust frameworks.

Phase 3: The *Advanced Plus* vault pattern is a recovery solution that introduces data security tools that scan and detect anomalies in the data entering and leaving the data vault. This forensic scanning of the data aims to improve data validity in the data vault and can provide an early warning of data inconsistencies across an organization.

All phases allow for organizations to plan how to recover the services and infrastructure needed to run the critical business services in an environment that is separate from the current production environment.

Use case 2: Backup protection (compute, storage, and database)

AWS Backup makes backups available until they reach the expiration of their retention periods. If a user (including the AWS account root user) attempts to delete a backup or change the lifecycle properties in a locked vault, AWS Backup denies the operation. A backup vault is a container that stores and organizes backups.

AWS Backup Vault Lock is an optional feature of a backup vault, which can be helpful in giving AWS customers additional security and control over their backup vaults. When a lock is active (immediately) in *compliance mode* and the grace time is over, the vault configuration cannot be altered or deleted by a customer, account or data owner, or AWS. Each vault can have one vault lock in place.[36]

Use case 3: Logically air-gapped AWS Backup Vault (compute, storage, and database)

AWS Backup has introduced a logically air-gapped vault[37]. With this new capability, the immutable backup copies are locked by default and further protected through encryption using keys owned by AWS. Employing a KMS key owned by AWS Backup to encrypt recovery points helps customers protect against accidental or unwanted deletions of customer managed keys.

Furthermore, the logically air-gapped vault simplifies the process of sharing its backup data with other accounts for restore purposes. Using AWS Resource Access Manager (AWS RAM), customers can share the vault data with specific accounts, including cross organization, for faster direct restore. After the vault is shared, backups can be directly restored, removing the step of copying backups into the destination account beforehand. This reduces the operational overhead, time to recover from a data loss event, and cost of extra copies.

500.17 Notices to superintendent **Requirement: (a)(1)** Notice of cybersecurity incident. Each covered entity shall notify the superintendent… in no event later than 72 hours after determining that a cybersecurity incident has occurred at the covered entity, its affiliates, or a third-party service provider. **Compliance date:** December 1, 2023	**Customer responsibility** This is an action for customers to complete independently. Customers should review their policies and procedures to determine if changes are needed because of this revision.	Not applicable.

Document revisions

Date	Description
November 2023	Initial draft
July 2024	First publication

Notes

[1] Improved, automated vulnerability management for cloud workloads with Amazon Inspector

[2] Use Amazon Inspector to manage your build and deploy pipelines for containerized applications

[3] How to scan your AWS Lambda functions with Amazon Inspector

[4] Amazon Inspector now scans AWS Lambda functions for vulnerabilities in application package dependencies

[5] Centrally deploy patching operations across your AWS Organization using Systems Manager Quick Setup

[6] Automate vulnerability management and remediation in AWS using Amazon Inspector and AWS Systems Manager-Part 1

[7] Automate vulnerability management and remediation in AWS using Amazon Inspector and AWS Systems Manager-Part2

[8] Implementing up-to-date images with automated EC2 Image Builder pipelines

[9] Policies and permissions in IAM

[10] Best practices for AWS Organizations service control policies in a multi-account environment | AWS Security Blog

[11] 5 ways to modernize workforce identity with AWS

[12] IAM Identity Center User Guide

[13] Use tags to manage and secure access to additional types of IAM resources | AWS Security Blog

[14] Scale cross-account AWS KMS-encrypted Amazon S3 bucket access using ABAC | AWS Security Blog

[15] Set Up Your AWS Notifications in One Place | AWS News Blog

[16] How to receive alerts when your IAM configuration changes | AWS Security Blog (amazon.com)

[17] Building a cloud CMDB on AWS for consistent resource configuration in hybrid environments | AWS Cloud Operations and Migrations Blog (amazon.com)

[18] Managing Assets Across Cloud Providers and Maintains Compliance Using AWS Config | Deutsche Börse Case Study | AWS (amazon.com)

[19] How you can use Amazon GuardDuty to detect suspicious activity within your AWS account | AWS Security Blog

[20] Endpoint Security: What is it and why it's important - AWS (amazon.com)

[21] New for Amazon GuardDuty – Malware Protection for Amazon EBS Volumes | AWS News Blog

[22] Cloud Storage on AWS

[23] Encrypting Data-at-Rest and -in-Transit - Logical Separation on AWS

[24] Introduction to the cryptographic details of AWS KMS

[25] AWS KMS concepts

[26] Application Load Balancer now supports TLS 1.3 (amazon.com)

[27] Cloud storage in minutes with AWS Storage Gateway (updated) | AWS Storage Blog (amazon.com)

[28] AWS Backup | Centralized Cloud Backup | FAQs (amazon.com)

[29] The anatomy of ransomware event targeting data residing in Amazon S3 | AWS Security Blog

[30] Data Protection Reference Architectures with AWS Backup (awsstatic.com)

[31] AWS Backup | Centralized Cloud Backup | FAQs (amazon.com)

[32] Disaster Recovery (DR) Architecture on AWS, Part I: Strategies for Recovery in the Cloud | AWS Architecture Blog
Disaster Recovery (DR) Architecture on AWS, Part II: Backup and Restore with Rapid Recovery | AWS Architecture Blog
Disaster Recovery (DR) Architecture on AWS, Part III: Pilot Light and Warm Standby | AWS Architecture Blog
Disaster Recovery (DR) Architecture on AWS, Part IV: Multi-site Active/Active | AWS Architecture Blog

[33] Step 9. Test data recovery capabilities - AWS Prescriptive Guidance (amazon.com)

[34] How to perform non-disruptive tests with AWS Elastic Disaster Recovery | AWS Storage Blog (amazon.com)

[35] Air-gapping is a common technique used to isolate the data vault network from the production environment network and the internet, manage access to backups, and provide a form of resiliency in the event of a disruption or outage.

[36] Enhance the security posture of your backups with AWS Backup Vault Lock | AWS Storage Blog (amazon.com)

[37] Introducing AWS Backup logically air-gapped vault | AWS Storage Blog (amazon.com)