

NIST Cybersecurity Framework (CSF) 2.0

Aligning to the NIST CSF 2.0 in the AWS Cloud

First Published January 2019

Updated January 2025



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.

© 2025 Amazon Web Services, Inc. or its affiliates. All rights reserved.

Contents

- Abstract..... 4
- Intended audience..... 5
- Introduction 5
- Security benefits of adopting the NIST CSF..... 9
- NIST CSF implementation use cases 10
 - Healthcare..... 10
 - Financial services 11
 - Supply chain services 11
 - International adoption 11
- NIST CSF and AWS Best Practices..... 12
- NIST CSF 2.0: Core Functions 13
 - CSF Core Function: Govern 14
 - CSF Core Function: Identify 19
 - CSF Core Function: Protect..... 23
 - CSF Core Function: Detect..... 30
 - CSF Core Function: Respond..... 34
 - CSF Core Function: Recover..... 38
- AWS services alignment with the CSF..... 41
- Conclusion 42
- Contributors..... 43
- Document revisions..... 43
- Notes..... 43

Abstract

Governments, industry sectors, and organizations around the world are increasingly recognizing the NIST Cybersecurity Framework (CSF) as a recommended cybersecurity baseline to help improve the cybersecurity risk management and resilience of their systems. In February 2024, the NIST published an updated version of the CSF, 2.0. This new version resulted from the work started in February 2022, when the NIST collected information through a Request for Information (RFI) for evaluating and improving cybersecurity resources, including the NIST CSF that was initially published in February 2014.

This paper evaluates the NIST CSF 2.0 and the many AWS Cloud offerings that public and commercial sector customers can use to align to the NIST CSF 2.0 to improve their cybersecurity posture.

Intended audience

This document is intended for cybersecurity professionals, risk management officers or other organization-wide decision makers considering how to implement a new or improve an existing cybersecurity framework in their organization. For details on how to configure the AWS services identified in this document, contact your [AWS Solutions Architect](#).

Introduction

The NIST Framework for Improving Critical Infrastructure Cybersecurity (NIST Cybersecurity Framework, or CSF) was originally published in February 2014 in response to Presidential Executive Order 13636, “Improving Critical Infrastructure Cybersecurity,” which called for the development of a voluntary framework to help organizations improve the cybersecurity, risk management, and resilience of their systems. NIST conferred with a broad range of partners from government, industry, and academia for over a year to build a consensus-based set of sound guidelines and practices.

The Cybersecurity Enhancement Act of 2014 reinforced the legitimacy and authority of the CSF by codifying it and its voluntary adoption into law, until the Presidential Executive Order on “Strengthening the Cybersecurity of Federal Networks and Critical Infrastructure” signed on May 11, 2017, mandated the use of CSF for all U.S. federal entities.

In February 2022, the NIST published a Request for Information (RFI) named “Evaluating and Improving NIST Cybersecurity Resources: The Cybersecurity Framework and Cybersecurity Supply Chain Risk Management.” After a consultation period and [discussion draft](#) published in August 2023, the NIST finalized the new [updated version](#) of the CSF in February 2024.

[As the NIST states](#): “*The NIST Cybersecurity Framework (CSF) 2.0 [...] offers a taxonomy of high-level cybersecurity outcomes that can be used by any organization — regardless of its size, sector, or maturity — to better understand, assess, prioritize, and communicate its cybersecurity efforts. The CSF does not prescribe how outcomes should be achieved. Rather, it links to online resources that provide additional guidance on practices and controls that could be used to achieve those outcomes.*”

[The latest 2.0 version](#) of the framework is designed to address all types of organizations, from small schools and nonprofits to large enterprises and government agencies, regardless of their current cybersecurity maturity. This further reinforces the

core cybersecurity disciplines that make up the CSF, which have been widely recommended by both government and industry as a baseline for any organization to adopt.

The CSF now has an expanded scope that goes beyond just protecting critical infrastructure like hospitals and power plants, to cover organizations across many sectors. It also introduces a new focus on cybersecurity governance, emphasizing that this is a key enterprise risk that senior leadership must consider alongside other strategic risks like financial and reputational risks.

The framework's core functions have been now organized into six key areas: Govern (new), Identify, Protect, Detect, Respond, and Recover. Together, these functions provide a comprehensive lifecycle approach to managing cybersecurity risk.

The updated CSF recognizes that organizations will have varying levels of experience and needs when it comes to implementing cybersecurity tools and practices. To support both new and experienced users, the framework includes new implementation examples and [quick-start guides](#) tailored for specific user profiles, such as small businesses, enterprise risk managers, and organizations looking to secure their supply chains.

This version improves alignment with other cybersecurity frameworks and standards or individual information and communications technology (ICT) risk management programs (for example, from the NIST Risk Management Framework). It incorporates more robust privacy risk management guidance, noting that *“Cybersecurity risk management is essential for addressing privacy risks related to the loss of the confidentiality, integrity, and availability of individuals’ data”*. There is also an increased emphasis on managing cybersecurity risks in the supply chain, with recommendations for organizations to use the CSF to foster cybersecurity risk oversight and communications with stakeholders across supply chains.

The CSF’s core concepts and structure, including its key functions, have been internationally recognized in the past through their alignment with other standards (as discussed in the following example), demonstrating the framework's global relevance and effectiveness.

In Feb 2018, the International Standards Organization released [ISO/IEC 27103:2018 — Information technology— Security techniques -- Cybersecurity and ISO and IEC Standards](#). This technical report provides guidance for implementing a cybersecurity framework based on existing standards. In fact, ISO 27103 promotes the same concepts and best practices that are reflected in the NIST CSF; specifically, a framework focused on security outcomes that is organized around five functions (Identify, Protect, Detect, Respond, Recover) and foundational activities mapped to existing standards, accreditations, and frameworks. Adopting this approach can help organizations achieve security outcomes while benefiting from the efficiencies of re-using instead of re-doing.

Figure 1 shows the *Cybersecurity standards usage for control systems in organizations worldwide in 2021*, as documented by the SANS Institute.¹ This diagram shows, through percentages, how frequently cybersecurity frameworks and standards are adopted (in a survey of 480 respondents).




Figure 1: Cybersecurity standards usage for control systems in organizations worldwide in 2021. Source: [Mark Bristow, SANS 2021 Survey: OT/ICS Cybersecurity](#)

According to the SANS Institute, 47.8% of 480 respondents indicated mapping their control systems to the NIST CSF standard. Using frameworks and standards helps organizations create defense control systems to support their cybersecurity efforts. According to Gartner, 16 US critical infrastructure sectors use the CSF and over 21 states have implemented it.² In addition to critical infrastructure and other private-sector organizations, other countries, including Italy and Israel, are using the CSF as the foundation for their national cybersecurity guidelines.

Since Fiscal Year 2016, U.S. Federal Information Security Modernization Act (FISMA) metrics have been organized around the CSF, and now reference it as a “standard for managing and reducing cybersecurity risks.” According to the FY16 FISMA Report to Congress, the Council of the Inspectors General on Integrity and Efficiency (CIGIE) aligned IG metrics with the five CSF functions to evaluate agency performance and promote consistent and comparable metrics and criteria between Chief Information Officer (CIO) and Inspector General (IG) assessments.

The most common applications of the CSF have manifested in three distinct scenarios:

- Evaluation of an organization’s enterprise-wide cybersecurity posture and maturity by conducting an assessment against the CSF model (Current Profile) determine the desired cybersecurity posture (Target Profile), and plan and prioritize resources and efforts to achieve the Target Profile.
- Evaluation of current and proposed products and services to meet security objectives aligned to CSF categories and subcategories to identify capability gaps and opportunities to reduce overlap/duplicative capabilities for efficiency.
- A reference for restructuring their security teams, processes, and training.

This paper identifies the key capabilities of AWS service offerings available globally that U.S. federal, state, and local agencies; global critical infrastructure owners and operators; as well as global commercial enterprises can leverage to align to the CSF (security **in** the cloud). It also provides support to establish the alignment of AWS Cloud services to the CSF (security **of** the cloud).

This means that you can use AWS services to deliver on the security objectives and outcomes identified in the CSF and that you can use AWS solutions to support your own alignment with the CSF and other required compliance standards. For U.S. federal agencies, in particular, using AWS solutions can facilitate your compliance with FISMA reporting metrics. This combination of outcomes will empower you with confidence in the security and resiliency of your data as you migrate critical workloads to the AWS Cloud.

Security benefits of adopting the NIST CSF

The CSF 2.0 offers a simple yet effective construct that consists of three elements—Core, Tiers, and Profiles. The *Core* represents a set of cybersecurity practices, outcomes, and technical, operational, and managerial security controls (referred to as Informative References) that support the six risk management functions—Govern, Identify, Protect, Detect, Respond, and Recover. The *Tiers* characterize an organization’s aptitude and maturity for managing the CSF functions and controls. NIST

CSF 2.0 considers four tiers: Partial, Risk Informed, Repeatable, and Adaptive. *Profiles* are intended to convey the organization’s “as is” and “to be” cybersecurity postures. Together, these three elements enable organizations to prioritize and address cybersecurity risks that are consistent with their business and mission needs.

It is important to note that implementation of the Core, Tiers, and Profiles is the responsibility of the organization that adopts the CSF (for example, government agency, financial institution, commercial start-up, and so on). This paper focuses on AWS services, solutions and capabilities supporting the Core that can enable you to achieve the security outcomes (Subcategories) in the CSF.

The Core references security controls from widely-adopted, internationally-recognized standards such as the [Cloud Controls Matrix \(CCM\) Version 4.0](#) (submitted by Cloud Security Alliance (CSA)), [NIST SP 800-53 Rev. 5](#), [CIS Controls](#) (submitted by Center for Internet Security (CIS)), and [CRI Profile Version 2.0](#) (submitted by Cyber Risk Institute). Customers can also find additional [informative references](#) on the NIST CSF 2.0 website.

While this list represents some of the most widely reputed standards, the CSF encourages organizations to use the controls catalog that best meets their organizational needs. The CSF was designed to be size-, sector-, and country-agnostic; therefore, public and private sector organizations should be assured that the CSF is applicable regardless of the type of organization or nation-state location.

NIST CSF implementation use cases

Healthcare

The U.S. Department of Health and Human Services completed a mapping of the Health Insurance Portability and Accountability Act of 1996 (HIPAA)³ Security Rule to the NIST CSF. Under HIPAA, covered entities and business associates must comply with the HIPAA Security Rule to ensure the confidentiality, integrity and availability of protected health information.⁴ Since HIPAA does not have a set of controls that can be assessed or a formal accreditation process, covered entities and business associates, like AWS, are HIPAA-eligible based on alignment with NIST 800-53- security controls that can be tested and verified in order to place services on the HIPAA eligibility list. The mapping between the NIST CSF and the HIPAA Security Rule promotes an additional layer of security since assessments performed for certain categories of the NIST CSF

may be more specific and detailed than those performed for the corresponding HIPAA Security Rule requirement.

Financial services

The U.S. Financial Services Sector Coordinating Council⁵ (FS-SCC) comprised of 70 financial services associations, institutions and utilities/exchanges, developed a sector-specific profile—a customized version of the NIST CSF that addresses unique aspects of the sector and its regulatory requirements.

The Financial Services Sector Specific Cybersecurity profile, drafted collaboratively with regulatory agencies, is a means to harmonize cybersecurity-related regulatory requirements. For example, the FS-SCC mapped the “Risk Management Strategy” category to nine different regulatory requirements and determined that the language and definitions, while different, largely addressed the same security objective.

Supply chain services

Supply chain services, such as Amazon Supply Chain Services (ASCS), implement the NIST CSF 2.0 to give merchants and sellers assurance that supply chain providers have taken reasonable measures to protect data and fortify technology involved in supply chain operations. The ASCS offering aligns cybersecurity controls with the NIST CSF 2.0 to manage information-security risk. This alignment demonstrates our overall commitment to upholding risk-based, outcome-focused security measures, including a defense-in-depth strategy, security-by-design approach, and continuous improvement of processes.

International adoption

Outside of the U.S., many countries have leveraged the NIST CSF for commercial and public sector use. Italy was one of the first international adopters of the NIST CSF and developed a national cybersecurity strategy against the five functions. In June 2018, the UK aligned its Minimum Cyber Security Standard—mandatory for all government departments—to the five functions.

Additionally, Israel and Japan localized the NIST CSF into their respective languages with Israel creating a cyber defense methodology based on its own adaptation of the NIST CSF. Uruguay performed a mapping of the CSF to ISO standards to strengthen connections to international frameworks. Switzerland, Scotland, Ireland, and Bermuda are also among the list of countries that are using the NIST CSF to improve cybersecurity and resiliency across their public and commercial sector organizations.

NIST provides a [list](#) of the different ways various countries around the world have adopted and adapted their national cybersecurity frameworks by using the NIST CSF 1.1. Many of these countries are now working to integrate their national frameworks with the new version, NIST CSF 2.0. As [stated](#) by NIST itself, “the CSF is used widely internationally; versions 1.1 and 1.0 have been translated into 13 languages, and NIST expects that CSF 2.0 also will be translated by volunteers around the world.”

In fact, CSF 2.0 translations have already been published directly on the NIST [website](#) with the aim of facilitating its adoption (currently available in Spanish, Portuguese, Polish, and Korean).

NIST CSF and AWS Best Practices

While this paper serves as a resource to provide organizational lifecycle risk management that connects business and mission objectives to cybersecurity activities, AWS also provides other best practices resources for customers moving their organizations to the cloud (AWS Cloud Adoption Framework) and customers designing, building or optimizing solutions on AWS (Well-Architected Framework).⁶

These resources supply complementary tools to support an organization in building and maturing their cybersecurity risk management programs, processes and practices in the cloud. More specifically, this NIST CSF whitepaper can be used in parallel with either of these best practices guides, serving as the foundation for your security program with Cloud Adoption Framework or Well-Architected Framework as an overlay for operationalizing the CSF security outcomes in the cloud.

For customers migrating to the cloud, the [AWS Cloud Adoption Framework \(AWS CAF\)](#) provides guidance that supports each unit in your organization so that each area understands how to update skills, adapt existing processes, and introduce new processes to take maximum advantage of the services provided by cloud computing.

Thousands of organizations around the world have successfully migrated their businesses to the cloud, relying on the AWS CAF to guide their efforts. AWS and our Partners provide tools and services that can help you every step of the way to ensure complete understanding and transition.

NIST CSF 2.0: Core Functions

The following table lists all Functions and Categories (each with their corresponding Identifier) that are part of the NIST CSF 2.0. We will analyze these one by one in the following sections of this whitepaper.

Function	Category	Category Identifier
Govern (GV)	Organizational Context	GV.OC
	Risk Management Strategy	GV.RM
	Roles, Responsibilities, and Authorities	GV.RR
	Policy	GV.PO
	Oversight	GV.OV
	Cybersecurity Supply Chain Risk Management	GV.SC
Identify (ID)	Asset Management	ID.AM
	Risk Assessment	ID.RA
	Improvement	ID.IM
Protect (PR)	Identity Management, Authentication, and Access Control	PR.AA
	Awareness and Training	PR.AT
	Data Security	PR.DS
	Platform Security	PR.PS
	Technology Infrastructure Resilience	PR.IR
Detect (DE)	Continuous Monitoring	DE.CM
	Adverse Event Analysis	DE.AE
Respond (RS)	Incident Management	RS.MA
	Incident Analysis	RS.AN
	Incident Response Reporting and Communication	RS.CO
	Incident Mitigation	RS.MI
Recover (RC)	Incident Recovery Plan Execution	RC.RP
	Incident Recovery Communication	RC.CO

CSF Core Function: Govern

This section addresses the six categories that comprise the Govern Function: Organizational Context; Risk Management Strategy; Roles, Responsibilities, and Authorities; Policy; Oversight; and Cybersecurity Supply Chain Risk Management.

The Govern Function is a new element in version 2.0 of the framework compared to the previous one. In version 1.0, some of the controls that are now grouped in the Govern Category were distributed in the other Categories. This update emphasizes the procedural and organizational activities that have an impact on the management of cyber security risk within organizations.

The Govern Function emphasizes the importance for organizations to have cybersecurity governance, starting from organizational context and including people, with their roles and responsibilities, in the cybersecurity strategy.

Following are the CSF core Categories for the Govern Function:

- **Organizational Context (GV.OC)** – The circumstances—mission, stakeholder expectations, dependencies, and legal, regulatory, and contractual requirements—surrounding the organization’s cybersecurity risk management decisions are understood.
- **Risk Management Strategy (GV.RM)** – The organization’s priorities, constraints, risk tolerance and appetite statements, and assumptions are established, communicated, and used to support operational risk decisions.
- **Roles, Responsibilities, and Authorities (GV.RR)** – Cybersecurity roles, responsibilities, and authorities to foster accountability, performance assessment, and continuous improvement are established and communicated.
- **Policy (GV.PO)** – Organizational cybersecurity policy is established, communicated, and enforced.
- **Oversight (GV.OV)** – Results of organization-wide cybersecurity risk management activities and performance are used to inform, improve, and adjust the risk management strategy.
- **Cybersecurity Supply Chain Risk Management (GV.SC)** – Cyber supply chain risk management processes are identified, established, managed, monitored, and improved by organizational stakeholders.

Customer responsibility

Governance is the “unsung hero” of cybersecurity. It lays the foundation and sets the standard for people, processes, and technology. AWS provides services and capabilities such as [AWS IAM Identity Center](#), [AWS Organizations](#), [AWS Control Tower](#), [AWS Config](#), [AWS Systems Manager](#), [AWS Service Catalog](#), [AWS Security Hub](#), [AWS](#)

[CloudTrail](#), and [Amazon CloudWatch](#) that customers can use to implement, monitor, and enforce governance. In this section, we'll discuss how you can use these and other AWS services for your governance goals.

For enterprise organizations, building a well-architected landing zone that is a multi-account AWS environment based on security and compliance best practices is the foundation for governing the development and growth of the organization without losing control of resources. AWS customers can use the [AWS Security Reference Architecture \(AWS SRA\)](#) as a holistic set of guidelines for deploying the full complement of AWS security services in a multi-account environment.

[AWS Organizations](#) is the service to start with for setting up a multi-account environment. Users and roles defined in the management account (which administers the entire organization) should be used only for tasks that must be performed by that account. In the management account, customers can define service control policies (SCPs), a type of organization policy that is used to manage permissions in the organization. SCPs are a security feature at the organization level and do not restrict users or roles in the management account, but they apply in the account organization's members (**GV.PO**). For this reason, workloads and resources should be stored in other member accounts in the organization and kept out of the management account.

Delegating policy management for the organizations to specified member accounts, to perform policy actions that are by default available only to the management account, is a best practice. Customers can also delegate a member account to administrate (provision, de-provision, and update roles across the organization's member accounts) and manage the [AWS IAM Identity Center](#) (the AWS solution for creating, managing, and connecting workforce users to AWS resources) from outside the management account, even though the IAM Identity Center instance must always reside in the management account of the organization. In this way, the number of people who require access to the management account is minimized, to help mitigate security concerns (**GV.RR**). In an organization in AWS Organizations, customers can organize AWS accounts according to a tree-structure, where each node of the tree, representing an organizational unit (OU), groups AWS accounts with common policies. In this way, policies can be defined at the OU level and inherited by the OU's children nodes.

[AWS Control Tower](#) automates the setup of a new landing zone based on the AWS Organizations account structure, using best-practices blueprints for identity, federated access, use of AWS Identity Center, and account structure.

Examples of blueprints that are automatically implemented in the landing zone include centralization of logging from [AWS CloudTrail](#) stored in [Amazon Simple Storage Service \(Amazon S3\)](#), with the possibility to investigate logs and create incident response mechanisms in integration with [Amazon EventBridge](#) and [Amazon CloudWatch](#).

Customers can define and automate the deployment of approved resources, obtaining a standardized landscape for product provisioning for their entire organization, by using [AWS Service Catalog](#). This service allows organizations to centrally manage commonly deployed IT services and helps organizations achieve consistent governance and meet compliance requirements. End users can quickly deploy only the approved IT services they need, following the constraints set by their organization (**GV.OV**).

The landing zone set up by AWS Control Tower is managed using a set of mandatory and optional controls. Mandatory controls are applied on your behalf by AWS Control Tower, while optional controls can be self-selected based on your unique needs so that accounts and configurations adhere to your policies. These controls refer to security standards such as the [CIS AWS Foundations Benchmark](#), [National Institute of Standards and Technology \(NIST\) SP 800-53 Rev. 5](#) and [Payment Card Industry Data Security Standard \(PCI DSS\)](#).

[AWS Config](#) is a service that enables customers to assess, audit, and evaluate the configurations of their AWS resources. Customers can enable AWS Config in each AWS account, but for a multi-account and multi-Region environment, resource data [aggregation](#) is a good option. An aggregator is an AWS Config resource type that can collect AWS Config configuration and compliance data from the following: 1) multiple accounts and multiple AWS Regions, 2) a single account and multiple AWS Regions, or 3) an organization in AWS Organizations and the accounts in that organization that have AWS Config enabled.

[AWS Security Hub](#) is a cloud security posture management (CSPM) service that performs security best practice checks, aggregates alerts coming from other AWS services (including AWS Config) and third-party software, and enables automated remediation. Findings coming from Security Hub in the multi-account environment can be centralized in an AWS account that is designated for security and audit functions and then investigated and prioritized in order to implement a runbook for managing and preventing security incidents (**GV.SC**).

Security Hub standards and controls don't guarantee compliance with regulatory frameworks or audits. Rather, the controls provide a way to monitor the current state of your AWS accounts and resources.

Security Hub also receives findings from [AWS IAM Access Analyzer](#), which supports customers in achieving the least-privilege principle. IAM Access Analyzer centrally reviews and removes unused and external access across AWS accounts with continuous monitoring and automated management and refinement with security integration workflows that alert security teams appropriately (**GV.RR**). For unused permissions, IAM Access Analyzer reviews existing policies and recommends a refined version that is tailored to your access activity (**GV.PO - GV.SC**).

Customers can benefit from AWS support for over 143 standards such as FedRAMP⁸, ISO, and PCI DSS. For more information, see the [AWS Compliance page](#). AWS provides information about its risk and compliance program to enable customers to incorporate AWS controls into their governance framework. This information can assist customers in documenting a complete control and governance framework, with AWS included as an important part of that framework. Services such as [Amazon Inspector](#) identify technical vulnerabilities, while [Amazon GuardDuty](#) includes intelligent threat detection mechanisms that can be fed into a risk posture and management process.

Even cost monitoring services such as [AWS Cost Explorer](#) and [AWS Budgets](#) can be used as mechanisms to monitor and help prevent unauthorized access that involves abnormal resource usage. With AWS Budgets, for example, you can implement alerting mechanisms that notify when a resource usage threshold is reached, both in terms of costs and usage. [AWS Trusted Advisor](#) is a valuable tool to continuously monitor the posture of an AWS account: It continuously evaluates your AWS environment by using best practice checks across the categories of cloud cost optimization, performance, resilience, security, operational excellence, and service limits, and it recommends actions to remediate deviations from best practices.

The enhanced visibility that the cloud provides increases the accuracy of a customer's risk posture, which allows risk decisions to be made based on more substantial data **(GV.OV)**.

You can use Security Hub findings collected from AWS security services such as Amazon GuardDuty, Amazon Inspector, and IAM Access Analyzer, as well as security findings from over 50 Partner solutions, to build a security data lake. [Amazon Security Lake](#) automatically centralizes security data from AWS environments, SaaS providers, on-premises, and cloud sources into a purpose-built data lake that is stored in your account. With Security Lake, you can get a more complete understanding of security data across your entire organization, including different AWS Regions **(GV.SC)**.

Because an organization evolves over time, as does the regulatory environment in which the organization operates, governance processes must be constantly monitored. [AWS Audit Manager](#) helps continually audit AWS usage to simplify how customers manage risk and compliance with regulations and industry standards. Audit Manager automates evidence collection so customers can more easily assess whether their own policies, procedures, and activities—also known as controls—are operating effectively. With Audit Manager, the organization is prepared to audit because customers can build audit-ready reports with much less manual effort and use them to manage stakeholders' control reviews **(GV.SC)**. During audit processes, customers can make use of [AWS Artifact](#) for the security of the cloud and related certifications and attestations: AWS Artifact provides

on-demand access to security and compliance reports from AWS and independent software vendors (ISVs) who sell their products on AWS Marketplace.

AWS responsibility

The purpose of the Govern Core Function in the NIST CSF 2.0 is to establish, communicate, and monitor the organization's cybersecurity risk management strategy. According to the NIST CSF 2.0, governance activities are crucial for incorporating cybersecurity into a broader Enterprise Risk Management (ERM) strategy. The Govern Core Function addresses understanding the organizational context, establishing a cybersecurity strategy, managing supply chain risks, defining roles, responsibilities, and authorities, setting security policies, and overseeing the overall cybersecurity strategy. The management of the security of the cloud at AWS is based precisely on these activities that are addressed by the Govern Function.

The AWS organizational structure provides a framework for planning, implementing, and controlling business operations. The organizational structure assigns roles and responsibilities to provide for adequate staffing, efficiency of operations, and the segregation of duties. Management also has established authority and appropriate lines of reporting for key personnel. AWS uses a structured on-boarding process to assist new employees as they become familiar with Amazon tools, processes, systems, policies, and procedures.

AWS maintains a formal risk management program to identify, analyze, treat, and continuously monitor and report risks that affect our business objectives, regulatory requirements, and customers. The AWS Risk Management (ARM) program identifies risks, documents them in a risk register, and reports results to leadership at least semi-annually.

The AWS organization has defined structures and reporting lines with assigned authority and responsibilities to appropriately adhere to the requirements that are relevant to security, availability, confidentiality, and privacy. AWS has an established information security organization managed by the AWS Security team, which is led by the AWS Chief Information Security Officer (CISO). AWS Security team responsibilities are defined and allocated across the organization.

AWS implements formal documented policies and procedures that provide guidance for operations and information security within the organization and the supporting AWS environments. Policies address purpose, scope, roles, responsibilities, and management commitment. AWS policies are maintained in a centralized location that is accessible by employees. Policies are reviewed and approved by AWS leadership at least annually.

The AWS risk management and assessment approach defines the periodic process of maintaining and continuously improving risk management. Through the use of reviews, evaluations, and risk management techniques, the objectives related to information security and quality are met.

AWS establishes and maintains a third-party risk scoring methodology, including third-party vendor categorization and vendor risk tiering.

Governance and Risk Management controls are reviewed by an independent auditor during several audits, including the ISO 27001⁹, ISO 27017, Cloud Computing Compliance Controls Catalog (C5), HIPAA, Information Security Registered Assessors Program (IRAP), NIST 800-53 (FedRAMP and DOD), SOC 2 COMMON CRITERIA, and FedRAMP audits.

CSF Core Function: Identify

This section addresses the three Categories that comprise the “Identify” Function: Asset Management, Risk Assessment, and Improvement (new) are intended to ensure that “the organization's current cybersecurity risks are understood”.

The four Categories Business Environment, Governance, Risk Management Strategy, and Supply Chain Risk Management have been withdrawn from this Function and have been incorporated into the newly created “Govern” Function.

Following are the CSF core Categories for Identify:

- **Asset Management (ID.AM)** – Assets (for example, data, hardware, software, systems, facilities, services, people) that enable the organization to achieve business purposes are identified and managed consistent with their relative importance to organizational objectives and the organization’s risk strategy.
- **Risk Assessment (ID.RA)** – The cybersecurity risk to the organization, assets, and individuals is understood by the organization.
- **Improvement (ID.IM)** – Improvements to organizational cybersecurity risk management processes, procedures, and activities are identified across all CSF Functions.

Customer responsibility

Identifying and managing IT assets is the first step in effective IT governance and security, and yet has been one of the most challenging. The Center for Internet Security (CIS)⁷ recognized the foundational importance of asset inventory and assigned physical and logical asset inventory as controls #1 and #2 of their Top 20. However, an accurate

IT inventory, both of physical assets and logical assets, is difficult to achieve and maintain for organizations of all sizes and resources.

Inventory solutions are limited in being able to identify and report on all IT assets across the organization for various reasons, such as network segmentation preventing the solution from “seeing” and reporting from various parts of the enterprise network, endpoint software agents not being fully deployed or functional, and incompatibility across a broad range of disparate technologies. Unfortunately, those assets that are “lost” or unaccounted for pose the greatest risk. If they are not tracked, they are most likely not receiving the most recent patches and updates, are not replaced during lifecycle refreshments, and malware may be allowed to exploit and maintain its hold of the asset.

Migrating to AWS provides two key benefits that can mitigate the challenges with maintaining asset inventories in an on-prem environment. First, AWS assumes sole responsibility for managing physical assets that comprise the AWS Cloud infrastructure. This can significantly reduce the burden of physical asset management for customers for those workloads that are hosted in AWS. The customer is still responsible for maintaining physical asset inventories for the equipment they keep in their environment (data centers, offices, deployed IoT, mobile workforce, and so on).

The second benefit is the ability to achieve deep visibility and asset inventory for logical assets hosted in a customer’s AWS account. This may sound like a bold claim, but it becomes quickly evident as it does not matter if an EC2 instance (virtual server) is turned on or off, whether the endpoint agent is installed and running, regardless of what network segment the asset is on, or other factors.

Whether using the AWS Management Console as a visual point-and-click interface, through the command line interface (CLI), or application programmable interface (API), customers can query and obtain visibility of AWS service assets. This reduces the inventory burden on the customer to the software they install on their EC2 instances and what data assets they store in AWS. AWS also has services that can perform this capability, like [Amazon Macie](#), which can identify, classify, label, and apply rules to data stored in [Amazon Simple Storage Service](#) (Amazon S3). You can also use [AWS Resource Explorer](#), a managed capability that makes it quicker to search and discover your AWS resources across AWS Regions and accounts with no additional costs.

You can use [AWS Systems Manager](#) to identify applications, servers, and other resources across AWS, on-premises, and other cloud environments through a unified experience. In particular, you can use [AWS Systems Manager Inventory](#) to collect *metadata* from your managed nodes. You can store, track, and determine which nodes are running the software and configurations required by your software policy, and which nodes need to be updated. You can also use [AWS Application Discovery Service](#)

specifically to collect and present data about the configuration, usage, and behavior of servers in your IT environments. From a people perspective, [AWS Identity and Access Management \(IAM\)](#) and [IAM Identity Center](#) help you get an accurate inventory of people that can access AWS accounts in your organization and the permissions granted to them. [AWS IAM Access Analyzer](#) helps you identify the resources in your organization and accounts, such as Amazon S3 buckets or IAM roles, that are shared with an external entity. This lets you identify unintended access to your resources and data, which represents a security risk. IAM Access Analyzer identifies resources shared with external principals by using logic-based reasoning to analyze the resource-based policies in your AWS environment. **(ID.AM)**.

AWS provides information about its risk and compliance program to enable customers to incorporate AWS controls into their governance framework. This information can assist customers in documenting a complete control and governance framework with AWS included as an important part of that framework. Services such as [Amazon Inspector](#) identify technical vulnerabilities that can be fed into a risk posture and management process. [AWS Systems Manager Patch Manager](#) helps manage patching across your Amazon EC2 instances. Amazon Inspector and Systems Manager findings can be viewed in [AWS Security Hub](#), a cloud security posture management service that helps automate AWS security checks and centralize security alerts.

[Amazon CodeGuru](#) is a static application security tool that uses machine learning to detect security policy violations and vulnerabilities. It provides suggestions for addressing security risks and generates metrics so you can track the security posture of your applications. CodeGuru Security's policies, which are informed by years of Amazon.com and AWS security best practices, help you to create and deploy secure, high-quality applications. [Amazon Q Developer](#) uses CodeGuru security detectors to scan your codebase for security vulnerabilities and code quality issues to improve the posture of your applications throughout the development cycle. You can initiate a scan of an entire codebase, analyzing each file in your local project or workspace, or enable auto scans that assess your code as you write it. When Amazon Q discovers a potential security vulnerability or quality issue in your code, it generates a finding with a description of the issue and a recommended fix. Some findings include an automatic fix, which updates your code files in-place. The enhanced visibility that the cloud provides increases the accuracy of a customer's risk posture, allowing risk-decisions to be made on more substantial data **(ID.RA)**.

You can use a combination of AWS Security Hub and [AWS Trusted Advisor](#) to identify areas for improving your security posture and receive recommendations to help follow AWS best practices **(ID.IM)**. You can find more details about these services in the CSF Core Function: Govern section.

AWS responsibility

AWS maintains stringent access control management by only providing data center access and information to employees and contractors who have a legitimate business need for such privileges. When an employee no longer has a business need for these privileges, their access is immediately revoked, even if they continue to be an employee of Amazon or AWS. Physical access to data centers by AWS employees is routinely logged and audited. Controls in place limit access to systems and data and provide that access to systems or data is restricted and monitored. In addition, customer data and server instances are logically isolated from other customers by default.

Privileged user access control is reviewed by an independent auditor during the AWS SOC 1, ISO 27001, PCI, and FedRAMP audits.

AWS risk management activities include the system development lifecycle (SDLC), which incorporates industry best practices and formal design reviews by the AWS Security Team, threat modeling and completion of a risk assessment. In addition, the AWS control environment is subject to regular internal and external risk assessments. AWS engages with external certifying bodies and independent auditors to review and test the AWS overall control environment.

AWS management has developed a strategic business plan which includes risk identification and the implementation of controls to mitigate or manage risks. AWS management re-evaluates the strategic business plan at least biannually. This process requires management to identify risks within its areas of responsibility and to implement appropriate measures designed to address those risks. In addition, the AWS control environment is subject to various internal and external risk assessments.

AWS Compliance and Security teams have established an information security framework and policies based on the Control Objectives for Information and related Technology (COBIT) framework and have effectively integrated the ISO 27001 certifiable framework based on ISO 27002 controls, American Institute of Certified Public Accountants (AICPA) Trust Services Principles, the PCI DSS v3.2, and the National Institute of Standards and Technology (NIST) Publication 800-53 Rev 4 (Recommended Security Controls for Federal Information Systems). AWS maintains the security policy, provides security training to employees, and performs application security reviews. These reviews assess the confidentiality, integrity, and availability of data, as well as alignment with the information security policy.

AWS Security regularly scans internet-facing service endpoint IP addresses for vulnerabilities (these scans do not include customer instances). AWS Security notifies the appropriate parties to remediate identified vulnerabilities. In addition, external

vulnerability threat assessments are performed regularly by independent security firms. Findings and recommendations resulting from these assessments are categorized and delivered to AWS leadership. These scans are done in a manner for the health and viability of the underlying AWS infrastructure and are not meant to replace the customer's own vulnerability scans required to meet their specific compliance requirements.

AWS maintains formal agreements with key third-party suppliers and implements appropriate relationship management mechanisms in line with their relationship to the business. The AWS third-party management processes are reviewed by independent auditors as part of AWS ongoing compliance with SOC and ISO 27001. In alignment with ISO 27001 standards, AWS hardware assets are assigned an owner, tracked and monitored by the AWS personnel with AWS proprietary inventory management tools. AWS procurement and supply chain team maintain relationships with all AWS suppliers. Refer to ISO 27001 standards; Annex A, domain 8 for additional details. AWS has been validated and certified by an independent auditor to confirm alignment with the ISO 27001 certification standard.

CSF Core Function: Protect

This section addresses the five Categories that comprise the "Protect" Function: Access Control, Awareness and Training, Data Security, Platform Security, and Technology Infrastructure Resilience.

The three Categories Information Protection, Processes, and Procedures; Maintenance; and Protective Technology have been withdrawn from this Function and have been incorporated into other Functions and Categories. The section also highlights AWS solutions that you can use to align to this Function.

The following are the CSF core Categories for Protect:

- **Identity Management, Authentication, and Access Control (PR.AA)** – Access to physical and logical assets is limited to authorized users, services, and hardware and managed commensurate with the assessed risk of unauthorized access.
- **Awareness and Training (PR.AT)** – The organization's personnel are provided with cybersecurity awareness and training so that they can perform their cybersecurity-related tasks.
- **Data Security (PR.DS)** – Data is managed consistent with the organization's risk strategy to protect the confidentiality, integrity, and availability of information.

- **Platform Security (PR.PS)** – The hardware, software (for example, firmware, operating systems, applications), and services of physical and virtual platforms are managed consistent with the organization's risk strategy to protect their confidentiality, integrity, and availability.
- **Technology Infrastructure Resilience (PR.IR)** – Security architectures are managed with the organization's risk strategy to protect asset confidentiality, integrity, and availability, and organizational resilience.

Customer responsibility

When looking at meeting the three security objectives of Confidentiality, Integrity, and Availability, the third can be very difficult to achieve in an on-premises environment with only one or two data centers. This is one of the greatest benefits of enterprise-scale cloud service providers, and AWS in particular, due to the AWS unique global infrastructure architecture (**PR.PS**).

You can distribute your application across multiple Availability Zones (AZs), which are logical fault isolation zones, each consisting of one or more discrete data centers within a Region. If architected properly with enhanced capacity management and automatic scaling capabilities, your application and data would not be impacted by the availability of a single data center. If you take advantage of all of the AZs in a Region (where each Region has three or more AZs), the loss of two data centers may still not have an impact to your application. Likewise, services such as Amazon S3 automatically replicate your data to at least three Availability Zones in the Region for a provided availability of 99.99% and data durability of 99.999999999% (**PR.IR**).

You can design your systems for confidentiality through encryption at rest and encryption in transit by using AWS encryption services such as [Amazon Elastic Block Store \(EBS\) encryption](#), [Amazon S3 encryption](#), Transparent Data Encryption for [Amazon Relational Database Service \(Amazon RDS\) SQL Server](#) and [RDS Oracle](#), and [VPN Gateway](#), or using your existing encryption solution. AWS supports TLS/SSL encryption (including TLS 1.3) for its API endpoints and the ability to create VPN tunnels to protect data in transit. AWS also provides the [AWS Key Management Service \(AWS KMS\)](#) and dedicated [Hardware Security Module](#) appliances to encrypt data at rest. You can choose to secure your data using the AWS provided capabilities, or use your own security tools.

Integrity can be facilitated by a variety of means. [Amazon CloudWatch](#) and [AWS CloudTrail](#) have integrity checks, customers can use digital signatures for API calls and logs, MD5 checksums can be employed in Amazon S3, and then there are numerous third-party solutions from our Partners. [AWS Config](#) helps to provide integrity of your AWS environment by monitoring for changes.

Within your AWS environment, AWS services such as [IAM](#), [Amazon Cognito](#), [IAM Identity Center](#), [AWS Directory Service](#), and features such as multi-factor authentication allow you to implement, manage, secure, monitor, and report on user identities, authentication standards, and access permissions **(PR.AA)**.

You are responsible for training your staff and end users on the policies and procedures for managing your environment. [AWS Training and Certification](#) includes digital and classroom training programs built by AWS Cloud experts. Learners of all skill levels can access over 600 free online courses in nearly 20 different languages in [AWS Skill Builder](#), our online learning center. Learners can also purchase individual and team subscriptions in AWS Skill Builder to access interactive learning such as hands-on labs, role-based games, and AWS certification exam preparation. AWS Training and Certification can help you achieve optimal return on investment (ROI) by equipping your workforce with cloud computing skills. In addition to free and subscription-based digital training, AWS offers in-depth classroom training delivered virtually or in person. You can opt for publicly available training courses by AWS or AWS Training Partners, or private onsite or virtual training so you can learn where and when it is convenient for you **(PR.AT)**.¹⁰

AWS provides the technical, operational, and contractual measures needed to protect your data. With AWS, you manage the privacy controls of your data, control how your data is used, who has access to it, and how it is encrypted to protect the confidentiality, integrity, and availability of data, in line with your organization's risk strategy. Services such as [AWS Identity and Access Management \(IAM\)](#) help you to securely manage access to AWS services and resources. [AWS CloudTrail](#) and [Amazon Macie](#) enable compliance, detection, and auditing, while [AWS CloudHSM](#) and [AWS Key Management Service \(KMS\)](#) allow you to generate and manage encryption keys—AWS makes it convenient to encrypt data in transit and at rest by using keys either managed by AWS or fully managed by you. You can also bring your own keys that were generated and managed outside of AWS **(PR.DS)**.

Today, the public-key cryptographic schemes we use to protect data are based on mathematical problems of factoring, discrete logarithms, and elliptic-curve [cryptography](#). While large-scale quantum computers are not yet publicly available, the creation of a sufficiently powerful quantum computer would accelerate solving the schemes we use in public-key cryptographic algorithms to protect our data today. AWS is already at work, preparing for a post-quantum world. The AWS Post-Quantum team interfaces with the global cryptographic community by participating in international conferences, the open literature, and standards organizations, with a goal of leading the adoption of quantum-resistant cloud-scale cryptographic technology. We are participating in projects and working groups on quantum-resistant cryptography, including the [NIST National Cybersecurity Center of Excellence \(NCCoE\) Migration to Post-Quantum Cryptography](#)

[project](#) and the [Open Quantum Safe initiative](#). For example, the [NIST Post-Quantum Cryptography standardization effort](#) is a competition-like process to solicit, evaluate, and standardize one or more quantum-resistant public-key cryptographic algorithms. The new public-key cryptography standards will specify one or more additional digital signatures, public-key encryption, and key encapsulation mechanism (KEM) algorithms to augment Federal Information Processing Standard (FIPS) 186-4, Digital Signature Standard (DSS), as well as NIST Special Publication 800-56A Revision 3. AWS team members have contributed to the key encapsulation mechanisms [ML-KEM](#) and the signature algorithms [ML-DSA](#) and [SLH-DSA](#). Today, AWS KMS supports a [hybrid post-quantum key exchange option](#) for the Transport Layer Security (TLS) network encryption protocol. You can use this TLS option when you connect to AWS KMS API endpoints. We're offering this feature before post-quantum algorithms are standardized so you can begin testing the effect of these key exchange protocols on AWS KMS calls.

You can improve Platform Security (**PR.PS**) by addressing multiple layers of the architectural stack. Network and application protection services on AWS enable you to enforce fine-grained security policy at every network control point across your organization. As you build your network using [networking services](#) on AWS, you have flexible options for where and how you build your network architecture, from defining private subnets to public, internet-facing networks. AWS provides your network and application security teams with services that help address their particular protection needs and compliance requirements. AWS network and application protection services give you fine-grained protections at the host-, network-, and application-level boundaries. [Amazon Virtual Private Cloud \(Amazon VPC\) security groups](#) provide protections at the host level for resources in your AWS workloads. For example, an RDS database can have its own security group per network interface with rules to allow inbound and outbound traffic to specific server IP ranges. At the network level, [AWS Network Firewall](#) allows you to tightly control traffic to, from, and between your VPCs with capabilities such as stateful inspection, intrusion prevention, and web filtering. For web application protection, services like AWS WAF allow you to filter parts of the web request, such as IP addresses, HTTP headers, HTTP body, or URI strings, to block common exploit patterns, such as SQL injection or cross-site scripting. [AWS Shield](#) helps protect your networks and applications from distributed denial of service (DDoS) and offers managed detection and response to fend off targeted exploits. Only at AWS can you get central management and visibility of these network and application security services in one place through integration with [AWS Firewall Manager](#). Firewall Manager simplifies the task of bringing new applications and resources into compliance by enforcing a common set of security rules. Finally, in terms of DNS protections across your [Amazon Virtual Private Cloud \(Amazon VPC\)](#) VPCs, Amazon Route 53 Resolver DNS Firewall helps block potentially malicious DNS queries.

[Cloud resilience](#) refers to the ability of an application to resist or recover from disruptions, including those related to infrastructure, dependent services, misconfigurations, transient network issues, and load spikes. Cloud resilience also plays a critical role in an organization's broader [business resilience](#) strategy, including the ability to meet [digital sovereignty](#) requirements. Resilient applications are those built with high [availability](#)—the percentage of time the application is available for use—and also those with a [disaster recovery](#) or continuity of operations plan in place. While AWS goes to great lengths to provide availability and reliability of the cloud, you share responsibility for maintaining availability and reliability within the cloud. Some best practices we recommend for building highly resilient systems in the AWS Cloud include designing for failure, automating failover and recovery, testing your recovery procedures, and accessing resources and reference architectures (**PR.IR**).

If you host all your resources in a single AZ that is affected by a failure, none of these resources would be available. It is therefore a best practice to architect across multiple AZs in the same Region to achieve extremely high Recovery Time Objectives (RTOs), Recovery Point Objectives (RPOs), and service availability. [Amazon Application Recovery Controller](#) allows you to automate management and coordination of recovery for your applications across AZs or Regions. For mission-critical applications, it is a best practice to architect across Regions to handle the case of an entire Region failing—perhaps as a result of a large-scale event caused by natural calamities. By monitoring a system for key performance indicators (KPIs), for example by using [Amazon CloudWatch](#), you can trigger automation when a threshold is breached. These KPIs should be a measure of business value and not of the technical aspects of the service operation. This allows for automatic notification and tracking of failures and automated recovery processes that work around or repair the failure. With sophisticated automation, it is possible to anticipate and remediate failures before they occur. One of the benefits of the cloud is that you can test how your system fails, and you can validate your recovery procedures. You can use a test environment to simulate different failures or recreate scenarios that led to failures in the past. This exposes failure pathways that you can test and fix before a real failure scenario, reducing the risk of components that have not been tested before failing.

Resiliency requires you to provide your organization with adequate resource capacity. You can use [AWS Auto Scaling](#) to monitor your applications and automatically adjust capacity to maintain an optimal level to satisfy demand without over- or under-provisioning. If you specify scaling policies, then AWS Auto Scaling can launch or terminate instances as demand on your application increases or decreases.

[AWS Resilience Hub](#) allows you to define, test, and track the resilience of your applications to make sure that you are able to meet your recovery objectives. With [AWS Fault Injection Service](#), you can improve application performance, observability, and

resilience through controlled fault injection experiments. [AWS Elastic Disaster Recovery \(AWS DRS\)](#) minimizes downtime and data loss with fast, reliable recovery of on-premises and cloud-based applications by using affordable storage, minimal compute, and point-in-time recovery.

[AWS Health Dashboard](#) provides alerts and remediation guidance when AWS is experiencing events that may impact your resources. The dashboard displays relevant and timely information to help you manage events in progress, and provides proactive notification to help you plan for scheduled activities. Alerts are triggered by changes in the health of AWS resources, giving you visibility into events and guidance to help quickly diagnose and resolve issues. The dashboard provides forward-looking notifications, and you can set up alerts across multiple channels, including email and mobile, so you receive timely and relevant information to help plan for scheduled changes that may affect you. In the event of AWS hardware maintenance activities that may impact one of your [Amazon Elastic Compute Cloud \(Amazon EC2\)](#) instances, for example, you would receive an alert with information to help you plan for, and proactively address, issues associated with the upcoming change.

AWS responsibility

AWS employs the concept of least privilege, whereby employee access is granted based on business need and job responsibilities, providing temporary role-based access to only those resources and data that are required at that moment in time.

AWS provides physical data center access only to approved employees. Employees who need data center access must first apply for access and provide a valid business justification. These requests are granted based on the principle of least privilege, where requests must specify to which layer of the data center the individual needs access, and are time-bound. Requests are reviewed and approved by authorized personnel, and access is revoked after the requested time expires. Once granted admittance, individuals are restricted to the areas that are specified in their permissions.

Third-party access is requested by approved AWS employees, who must apply for third-party access and provide a valid business justification. These requests are granted based on the principle of least privilege, where requests must specify to which layer of the data center the individual needs access, and are time-bound. These requests are approved by authorized personnel, and access is revoked after the requested time expires.

Once granted admittance, individuals are restricted to the areas that are specified in their permissions. Anyone granted visitor badge access must present identification when arriving on site and are signed in and escorted by authorized staff.

Furthermore, AWS makes the [AWS Nitro System](#) available to customers. The Nitro System provides enhanced security that continuously monitors, protects, and verifies the instance hardware and firmware. Virtualization resources are offloaded to dedicated hardware and software, minimizing the threat landscape. Finally, the Nitro System security model is locked down and prohibits administrative access, reducing the possibility of human error and tampering.

In the last calendar quarter of 2022, AWS engaged NCC Group to conduct an architecture review of the AWS Nitro System design, with a focus on specific claims AWS made for the security of the Nitro System APIs. Following the NCC Group assessment activities, AWS Nitro System got [independent affirmation](#) of its confidential compute capabilities.

AWS has implemented formal, documented security awareness and training policies and procedures for our employees and contractors that address purpose, scope, roles, responsibilities, management commitment, coordination among organizational entities, and compliance.

AWS [FedRAMP](#) and [ISO 27001](#) certifications document in detail the policies and procedures by which AWS operates, maintains, controls, approves, deploys, reports, and monitors changes to its environment and infrastructure, as well as how AWS provides redundancy and emergency responses for its physical infrastructure.

Additionally, the certifications document in detail the manner in which remote maintenance for AWS services is approved, performed, logged and reviewed so as to prevent unauthorized access. They also address the manner in which AWS sanitizes media and destroys data. AWS uses products and procedures that align with NIST Special Publication 800-88 Guidelines for Media Sanitization. You are also responsible for preparing the policies, processes, and procedures for data protection.

To support billing and maintenance requirements, AWS assets are assigned an owner, tracked, and monitored with AWS proprietary inventory management tools. AWS asset owner maintenance procedures are carried out by using a proprietary tool with specified checks that must be completed according to the documented maintenance schedule. Third-party auditors test AWS asset management controls by validating that the asset owner is documented and that the condition of the assets is visually inspected according to the documented asset management policy.

AWS services can also greatly facilitate the task of managing and performing systems maintenance for our customers. First, based on AWS infrastructure previously discussed with regard to Availability Zones (AZs), an application that was architected for high availability across multiple AZs can allow you to segregate maintenance activities. You can take assets within an AZ offline for maintenance without affecting the performance of

the overall application, as the duplicate assets in the other AZs scale out and pick up the load.

Maintenance can be accomplished one AZ at a time, and can be automated with stop-gates and reporting as required. In addition, entire architectures can be shifted over from a dev/test (Blue) environment to an operations (Green) environment, and vice versa, where that method is desired.

CSF Core Function: Detect

This section addresses the two Categories that comprise the “Detect” Function: Continuous Monitoring and Adverse Event Analysis, which have been renamed from the CSF 1.1 version (respectively they were “Security Continuous Monitoring” and “Anomalies and Events”). The Detection Processes Category, previously part of the CSF 1.1, has been withdrawn from this Function and incorporated into other Categories and Functions (such as Govern, Identify, and Detect Adverse Event Analysis). This section also highlights AWS services that you can use to align to this Function.

The following are the CSF core Categories for Detect:

- **Continuous Monitoring (DE.CM)** – Assets are monitored to find anomalies, indicators of compromise, and other potentially adverse events.
- **Adverse Event Analysis (DE.AE)** – Anomalies, indicators of compromise, and other potentially adverse events are analyzed to characterize the events and detect cybersecurity incidents.

Customer responsibility

The ability to gather, analyze, and alert on security-relevant events is fundamental to any cybersecurity risk management program. The API-driven nature of cloud technology provides a new level of visibility and automation not previously possible. With every action taken resulting in one or more audit records, AWS provides a wealth of activity information available to customers within their account structure. However, the volume of data can present its own challenges.

Finding the proverbial “needle in the haystack” is a real problem, but the capacity and capabilities the cloud provides are well-suited to resolve these challenges. With the appropriate log processing infrastructure, automation, and data analysis, it is possible to achieve near-real-time detection and response for critical events while filtering out false-positives and low/accepted risks.

AWS offers services that can be utilized as part of a comprehensive Security Operations strategy for nearly continuous monitoring and threat detection. For

Continuous Monitoring (**DE.CM**), there is [AWS CloudTrail](#) for logging API calls, where the logs can be digitally signed and encrypted and then stored in a secure [Amazon S3](#) bucket. CloudTrail records important information about each action, including who made the request, the services used, and the actions performed, thus helping you prove compliance, improve security posture, and consolidate activity records across Regions and accounts. This maps to Subcategories such as **DE.CM-01: Networks and network services are monitored to find potentially adverse events** and **DE.CM-03: Personnel activity and technology usage are monitored to find potentially adverse events**.

For **DE.CM-01**, [Virtual Private Cloud \(VPC\) Flow Logs](#) monitor the network activity going in and out of your VPC. This enables you to gain operational visibility into your network dependencies and traffic patterns, detect anomalies and help prevent data leakage, or troubleshoot network connectivity and configuration issues, while also archive flow logs to help meet compliance requirements.

There is also [Amazon CloudWatch](#), a monitoring service for resources and applications running within or even outside AWS. You can use CloudWatch to collect and track metrics, collect and monitor log files, and set alarms, mapping to **DE.CM-09: Computing hardware and software, runtime environments, and their data are monitored to find potentially adverse events**. Amazon CloudWatch allows you to gain system-wide visibility into resource utilization, application performance, and operational health.

[Amazon GuardDuty](#) is an intelligent threat detection service, fully managed, with integrated threat intelligence, machine learning (ML) anomaly detection, and malware scanning. GuardDuty continuously monitors AWS accounts, workloads, runtime activity, and data (**DE.CM-09: Computing hardware and software, runtime environments, and their data are monitored to find potentially adverse events**). If potential malicious activity, such as anomalous behavior, credential exfiltration, or command and control infrastructure communication is detected, GuardDuty generates detailed security findings that are designed to be integrated with existing event management and workflow systems and can be used for security visibility and assisting in remediation (**DE.CM-09**).

Another security service for detection is [Amazon Inspector](#), an automated vulnerability management service that continually scans instances and code, delivering near real-time vulnerability findings. Amazon Inspector shows a highly contextualized and meaningful risk score for each finding, so that you can set more accurate response priorities.

The **DE.CM** Category does not apply only to hardware and software, but also data (**DE.CM-09**). [Amazon Macie](#) discovers sensitive data by using machine learning and pattern matching, provides visibility into data security risks, and enables automated protection against those risks. In Macie, a finding is a detailed report of a potential

policy violation related to sensitive data. Macie automatically detects a large and growing list of sensitive data types, including personally identifiable information (PII) such as names, addresses, and credit card numbers.

For a more holistic approach, customers want to gain visibility into their AWS accounts' security posture. [AWS Config](#) is a fully managed service that provides resource inventory, configuration history, and configuration change notifications for security and governance. These capabilities use compliance auditing, security analysis, resource change tracking, and troubleshooting.

AWS Config is the foundation for [AWS Security Hub](#), a cloud security posture management (CSPM) service that performs automated, continuous security best practice checks against AWS resources to help identify misconfigurations. Security Hub security checks make use of configuration items recorded by AWS Config and support controls for multiple industry best practices and standards, including [AWS Foundational Security Best Practices](#), [CIS AWS Foundations Benchmark](#) (with the recently added support for the 3.0.0 version) and [National Institute of Standards and Technology \(NIST\) SP 800-53 Rev. 5](#).

Finally, [AWS Health](#) is the authoritative data source for events and changes that affect AWS cloud resources. AWS Health notifies about service events, planned changes, and account notifications to help you manage and take actions.

In the next section, we are going to focus on the second Category: **Adverse Event Analysis (DE.AE)**, describing AWS services that can help analyze security-related findings to characterize and investigate events and detect cybersecurity incidents.

[AWS Security Hub](#) can also automatically aggregate security findings in a standardized data format from AWS and partner services, so that you can more easily enrich, investigate, and remediate them. You can apply and save filters to create contextual views by specific criteria, and prioritize the data and view of your organization's security posture that fits your needs, estimating findings' impact through their security scores (**DE.AE-02: Potentially adverse events are analyzed to better understand associated activities**, **DE.AE-03: Information is correlated from multiple sources**, **DE.AE-04: The estimated impact and scope of adverse events are understood**).

The same Subcategories bring us to [Amazon Security Lake](#), a service that automates the sourcing, aggregation, normalization, and data management of security data from AWS environments, SaaS providers, on-premises, and cloud sources into a security data lake. Security Lake facilitates solutions that perform streaming analytics detections, time-series analytics, user and entity behavior analytics (UEBA), security orchestration and remediation (SOAR), and incident response, all through the Open Cybersecurity Schema Framework (OCSF) format. It is also possible to query and

visualize data from Amazon Security Lake by using [Amazon Athena](#) and [Amazon QuickSight](#), performing ad hoc analysis and quickly getting business insights from security-related data.

[Amazon Detective](#) automatically collects log data from AWS resources and uses machine learning, statistical analysis, and graph theory to conduct faster and more efficient security investigations, quickly identifying the root cause of potential security issues or suspicious activities. Detective analyzes patterns and behaviors in the findings, such as similar types of unauthorized or suspicious activities, to determine relationships and group them accordingly. Findings that share similar Tactics, Techniques, and Procedures (TTPs), as described in frameworks like MITRE ATT&CK, are grouped together to highlight potential coordinated exploits (**DE.AE-02**, **DE.AE-03**, **DE.AE-07**: *Cyber threat intelligence and other contextual information are integrated into the analysis*).

Services like GuardDuty, Amazon Inspector and Macie can also help you estimate the impact and scope of adverse events (**DE.AE-04**: *The estimated impact and scope of adverse events are understood*). [AWS Shield](#) and [AWS WAF](#) even use advanced threat intelligence information, so that threats are mitigated early by initiating automated responses (**DE.AE-07**).

Finally, customers may even choose to take advantage of AWS analytics and artificial intelligence (AI) and machine learning (ML) services to model and analyze log data, such as [Amazon OpenSearch Service](#) and [Amazon SageMaker](#).

AWS responsibility

AWS has implemented a formal, documented incident response policy and program. The policy addresses purpose, scope, roles, responsibilities, and management commitment.

AWS contingency plans and incident response playbooks have defined and tested tools and processes to detect, mitigate, investigate, and assess security incidents. AWS security engineers follow a protocol when responding to potential data security incidents.

AWS uses a wide variety of automated monitoring systems to provide a high level of service performance and availability. AWS defines a security incident as a security-related adverse event in which there was a loss of data confidentiality, disruption of data or systems integrity, or disruption or denial of availability. AWS monitoring tools are implemented to detect unusual or unauthorized activities and conditions at ingress and egress communication points. These tools monitor server and network usage, port scanning activities, application usage, and unauthorized intrusion attempts. Staff operators in the US, EMEA (Europe, the Middle East, and Africa), and APAC (Asia-

Pacific) provide 24 x 7 x 365 coverage to detect incidents and to manage the impact and resolution.

AWS maintains the [AWS Security Bulletins webpage](#) to notify customers of security and privacy events that affect AWS services. Customers can subscribe to the Security Bulletin RSS feed to be notified of security announcements on the Security Bulletin webpage. The AWS Support team maintains a [Service Health Dashboard webpage](#) to alert customers to broadly impacting availability issues.

CSF Core Function: Respond

This section addresses the four categories that comprise the “Respond” Function: Incident Management, Incident Analysis (previously *Analysis*), Incident Response Reporting and Communication (previously *Communications*) and Incident Mitigation (previously *Mitigation*). Compared to the CSF 1.1 version, one Category—Improvements—has been incorporated into the Identify Function (**ID.IM**), whereas the other one—Response Planning—has become part of Incident Management, which represents a standalone Category in the CSF 2.0. This section will also summarize the key AWS services that can help you align to this Function.

Following are the CSF core Categories for Respond:

- **Incident Management (RS.MA)** – Responses to detected cybersecurity incidents are managed.
- **Incident Analysis (RS.AN)** – Investigations are conducted to ensure effective response and support forensics and recovery activities.
- **Incident Response Reporting and Communication (RS.CO)** – Response activities are coordinated with internal and external stakeholders as required by laws, regulations, or policies.
- **Incident Mitigation (RS.MI)** – Activities are performed to prevent expansion of an event and mitigate its effects.

Customer responsibility

The time between detection and response is critical. Well-run, repeatable response plans minimize exposure and accelerate recovery. Automation enabled by the cloud allows for the implementation of sophisticated playbooks as code with much quicker response times.

By simply tagging an [Amazon Elastic Compute Cloud](#) (Amazon EC2) instance, for example, automation can isolate the instance, take a forensic snapshot, install analysis tools, connect the suspect instance to a forensic workstation, and cut a ticket to a cybersecurity analyst. The capabilities listed below facilitate the creation of automated processes to add speed and consistency to your incident response processes.

Moreover, these tools allow you to maintain a history of the communications for use in a post-event review.

The first Category within this Function—**Incident Management (RS.MA)**—mostly revolves around incident management processes such as prioritization, incident categorization, reporting and validation, escalation, and incident plan execution. [AWS Security Incident Response](#) helps you respond when it matters most. The service combines the power of automated monitoring and investigation, accelerated communication and coordination, and direct 24/7 access to the AWS Customer Incident Response Team (CIRT) to quickly prepare for, respond to, and recover from security events. These features also make the service suitable to help you address other Categories, such as **Incident Analysis (RS.AN)** and **Incident Response Reporting and Communication (RS.CO)**.

Particularly for **Incident Analysis (RS.AN)**, since we have already talked about **Adverse Event Analysis (DE.AE)**, it's useful to remind ourselves what the difference is between a security event and a security incident. [NIST defines an event](#) as “Any observable occurrence in a network or information system,” whereas a [cybersecurity incident](#) is “A cybersecurity event that has been determined to have an impact on the organization prompting the need for response and recovery.” This means that, while a security incident usually stems from an event, the opposite is not necessarily true: A security-related event might not generate an incident. This brief clarification is necessary to introduce the RS.AN scope.

AWS services such as [AWS CloudTrail](#) and [Amazon Detective](#), other than offering features for adverse events analysis, can address the incident analysis topic, helping you understand what has taken place during an incident and the root cause of the incident (**RS.AN-03: Analysis is performed to establish what has taken place during an incident and the root cause of the incident**).

Let's move to the **Incident Mitigation (RS.MI)** Category. We have [AWS Network Firewall](#), a managed service that you can use to deploy essential network protections for virtual private clouds (VPCs). The Network Firewall flexible rules engine lets you define firewall rules that offer fine-grained control over network traffic, such as blocking outbound requests to prevent the spread of malicious activity. The rules engine is highly flexible, so you can build custom firewall rules to protect your unique workloads. While the service supports thousands of rules—and the rules can be based on domain, port, protocol, IP addresses, and pattern matching—you can also import rules you have already written in common open source rule formats or import compatible rules that are sourced from AWS partners.

Another key aspect of incidents containment revolves around permissions. Within [AWS Identity and Access Management \(AWS IAM\)](#), each IAM entity (users, groups, or roles) has its own IAM policy attached to it, making sure the entity can only perform actions

that are explicitly allowed through the combination of policies involved in that specific request. By combining this with the implementation of a strategy based around the least privilege principle, you can make sure that even if an entity gets compromised, each request made by it will still be evaluated against the IAM policy, thus limiting the scope of impact of the incident.

A similar concept applies to features such as [security groups](#) and [network access control lists \(network ACLs\)](#). Security groups can be used to help secure instances within a VPC, enabling customers to specify both inbound and outbound network traffic that is allowed to or from each Amazon EC2 instance (traffic which is not explicitly allowed is automatically denied). Network traffic entering and exiting each subnet can be allowed or denied by using network ACLs, which can be used to set both Allow and Deny rules and perform stateless filtering—in contrast to security groups, whose filtering is stateful. By controlling which traffic is allowed or denied, you can help prevent a compromised instance from being used for lateral movement.

[AWS Shield](#) is a managed service that provides protection against DDoS for applications running on AWS, with AWS Shield Advanced providing additional protections against more sophisticated and larger exploits for your applications. With [automatic application-layer DDoS mitigation](#), you can configure Shield Advanced to respond automatically to mitigate application layer (layer 7) attempts against your protected application layer resources, by counting or blocking web requests that are part of the attempt.

As for eradication (**RS.MI-02: Incidents are eradicated**), services such as [AWS Systems Manager](#), [AWS Elastic Disaster Recovery](#), and [AWS Backup](#) can help achieve the goals defined through this Subcategory. AWS Systems Manager is the operations hub for AWS applications and resources and an end-to-end management solution for hybrid and multicloud environments that facilitates secure operations at scale. One of its key capabilities, Patch Manager, automates the process of patching managed nodes with both security-related updates and other types of updates.

We will look deeper into disaster recovery and backup when discussing the Recover Function.

During the eradication phase, you can also use services such as [AWS CloudTrail](#), [IAM](#), [Amazon Inspector](#), and [Amazon GuardDuty](#) to identify and delete malware, disable breached user credentials, as well as mitigate vulnerabilities that were exploited. By combining these services together, you can remove potentially malicious software that was installed on your instances, bring your environment to its before-incident status, apply patches for both operating systems and applications, and reduce the chances of the same incident happening again.

While the cloud does offer capabilities to streamline and expedite the collection and dissemination of information, there is always a human element involved in response coordination. Cybersecurity analysis requires investigative action, forensics, and understanding of the incident. These necessarily demand some level of human interaction.

Even so, it is of utmost importance to think about ways to automate actions, whenever possible. By using automation, you can set up response and remediation activities more rapidly in a standardized and repeatable manner. Manual processes are error prone, unreliable, and time consuming, which is why the use of automation is critical to realizing the full benefits of the cloud. This can include activities for incident containment, eradication, and communication, such as isolating compromised instances, installing patches, and notifying security teams about a potential incident.

[AWS Security Hub](#) can also work as a starting point for automated response and remediation workflows through its integration with other AWS services, thus accelerating the mean time to resolution. As an example, you can create custom automated response, remediation, and enrichment workflows using the Security Hub integration with [Amazon EventBridge](#). You can also use other integrations to create Security Orchestration Automation and Response (SOAR) and security information and event management (SIEM) workflows.

Other than combining AWS services together, customers can use [AWS Solutions](#). Built by AWS and AWS Partners, these solutions bring together AWS services into pre-configured deployable code, architecture diagrams, and implementation guides, tailored to meet customer use cases. [Automated Security Response on AWS](#) is a solution that works with AWS Security Hub and provides predefined response and remediation actions based on industry compliance standards and best practices for security threats.

Finally, let's discuss AWS services that you can use for **Incident Response Reporting and Communication (RS.CO)**.

[Amazon EventBridge](#) is a serverless service that uses events to connect application components together, making it easier for you to build scalable event-driven applications. Event-driven architecture is a style of building loosely coupled software systems that work together by emitting and responding to events. Event-driven architectures can help automatically perform actions in response to events. For example, notification and communication actions can be implemented thanks to the integration between EventBridge and other services like [Amazon Simple Notification Service \(Amazon SNS\)](#). Amazon SNS is a web service that you can use to set up, operate, and send notifications from the cloud. It provides you with a highly scalable, flexible, and cost-effective capability to publish messages from an application and

immediately deliver them to subscribers or other applications, following the “publish-subscribe” (pub-sub) messaging paradigm.

AWS responsibility

AWS has implemented a formal, documented incident response policy and program. The policy addresses purpose, scope, roles, responsibilities, and management commitment.

AWS utilizes a three-phased approach to manage incidents:

- Activation and notification phase
- Recovery phase
- Reconstitution phase

To maintain the effectiveness of the AWS Incident Management plan, AWS conducts incident response testing. This testing provides excellent coverage for the discovery of previously unknown defects and failure modes. In addition, it allows the Amazon Security and Service teams to test the systems for potential customer impact and further prepare staff to handle incidents such as detection and analysis, containment, eradication, and recovery, and post-incident activities.

The Incident Response Test Plan is run annually, in conjunction with the Incident Response plan. AWS Incident Management planning, testing and test results are reviewed by third-party auditors.

CSF Core Function: Recover

This section addresses the two categories that comprise the “Recover” Function: Incident Recovery Plan Execution and Incident Recovery Communication. It also summarizes the key AWS solutions that you can leverage to align to this Function. The Recover Function refers to how assets and operations that are affected by a cybersecurity incident are restored.

Recovery strategies are a core function in the risk management of every organization: Any occurrence that actually or potentially jeopardizes the confidentiality, integrity, or availability of an information system can be defined as a security incident. Cybersecurity incidents, such as security misconfigurations and cyberattacks, can be considered a type of disaster event, for which an organization needs to plan a disaster recovery strategy. One of the benefits of using the AWS Cloud is the availability of a broad set of disaster recovery strategies, depending on the classification of the data and workload. We can classify them in four categories according to the established RPO (recovery point objective) and RTO (recovery time objective) that play a

fundamental role in planning disaster recovery processes. Customers should consider the Backup and Restore, Pilot Light, Warm Standby, or Multi-Site Active/Active strategy for their own applications, and also to align with regulatory requirements. (For more information on these strategies, see the [AWS Well-Architected Framework](#).)

Regularly assessing and testing the disaster recovery strategy is also a best practice for improving the organization's ability to restore the healthy state and activities prior to a security incident.

Following are the CSF core Categories for Recover:

- **Incident Recovery Plan Execution (RC.RP)** – Restoration activities are performed to provide operational availability of systems and services affected by cybersecurity incidents.
- **Incident Recovery Communication (RC.CO)** – Restoration activities are coordinated with internal and external parties.

Customer responsibility

Customers are responsible for planning, testing, and performing recovery operations for their applications and data to maintain their business continuity. The cause of an outage may come from many different sources. AWS services provide many advanced capabilities for self-healing and automated recovery.

For example, the use of Auto Scaling groups across multiple Availability Zones allows for the infrastructure to monitor the health of Amazon EC2 instances and rapidly replace a failed instance with a new Amazon Machine Image (AMI). Additionally, the use of Amazon CloudWatch, AWS Lambda, and other services/service capabilities can automate recovery actions to include everything from deploying an entire AWS environment and application, to failing over to a different AWS Region, restoring data from backups, and more.

Depending on the disaster recovery strategy, customers can use services such as [AWS Backup](#) to implement automatic and periodic backup processes, so that a backup of the data, preferably encrypted, is always available for the purpose of recovery. AWS Backup is a fully managed service that you can use to centralize and configure backup policies and monitor activity for AWS resources in one place. Backup and restore is a suitable approach for mitigating against data loss or corruption. Strategies for mitigating ransomware attempts, for example, involve the use of multiple AWS accounts, so that if an adversary were able to penetrate the AWS account on which the data resides, the backup of that data would reside on a different AWS account, making it necessary to cross an additional border to reach that data. Backup with point-in-time recovery is available by using snapshots in several services such as [Amazon Elastic Block Store](#), [Amazon DocumentDB](#), [Amazon Relational Database Service \(Amazon RDS\)](#), [Amazon](#)

[Aurora](#), [Amazon Elastic File System \(Amazon EFS\)](#), [Amazon Redshift](#), [Amazon Neptune](#), and [Amazon FSx](#). For data stored in [Amazon S3](#), customers can asynchronously copy objects to an S3 bucket in the disaster recovery Region continuously, while providing versioning for the stored objects so that they can choose a restoration point.

In addition to data, customers must redeploy the infrastructure, the configuration, and the application code in the recovery site, which can be another Availability Zone in the same AWS Region, another AWS Region in the same AWS account or a different AWS account. Infrastructure as code (IaC) that uses services such as [AWS CloudFormation](#) or the AWS Cloud Development Kit (AWS CDK) is a best practices to enable recovery without errors, reduce recovery time, and also track changes over time. In addition to user data, be sure to also back up code and configuration, including Amazon Machine Images (AMIs) you use to create Amazon EC2 instances. You can use AWS CodePipeline to automate redeployment of application code and configuration. To periodically evaluate their ability to manage incidents and simulate the recovery phase, customers can use the [AWS Resilience Hub](#), which uses IaC mechanisms. Within AWS Resilience Hub, customers can also create and run Amazon Fault Injection Service experiments, which mimic real-life disruptions to applications to help customer better understand dependencies and uncover potential weaknesses, including in the recovery of the applications. [Amazon Application Recovery Controller](#) (ARC) helps customers prepare for and complete faster recovery for applications running on AWS.

With the Pilot Light approach, customers replicate their own data from one AWS Region to another and provision a copy of core workload infrastructure. [AWS Elastic Disaster Recovery \(AWS DRS\)](#) continuously replicates server-hosted applications and server-hosted databases from a source into AWS using block-level replication of the underlying server. Elastic Disaster Recovery enables you to use a Region in the AWS Cloud as a disaster recovery target for a workload that is hosted on-premises or on another cloud provider. It can also be used for disaster recovery of AWS hosted workloads if they consist only of applications and databases that are hosted on EC2 (that is, not on Amazon RDS). Elastic Disaster Recovery uses the Pilot Light strategy, maintaining a copy of data and “switched-off” resources in a VPC used as a staging area. When a failover event is triggered, the staged resources are used to automatically create a full-capacity deployment in the target VPC used as the recovery location. **(RC.RP)**

Lastly, actions that involve public relations, reputation management, and communicating recovery activities are an important part of how organizations handle events that impact their environment. Services such as Amazon EventBridge and Amazon Simple Notification Service help customers in notification and communications both internal and external, sending events or emails to stakeholders that need to be informed of the event and to recover. **(RC.CO)**

AWS responsibility

The AWS Resiliency Program encompasses the processes and procedures by which AWS identifies, responds to, and recovers from a major availability event or incident within the AWS services environment. This program builds upon the traditional approach of addressing contingency management, which incorporates elements of business continuity and disaster recovery plans and expands this to consider critical elements of proactive risk mitigation strategies, such as engineering physically separate Availability Zones (AZs) and continuous infrastructure capacity planning.

The AWS business continuity plan details the three-phased approach that AWS has developed to recover and reconstitute the AWS infrastructure:

- Activation and notification phase
- Recovery phase
- Reconstitution phase

This approach allows AWS to perform system recovery and reconstitution efforts in a methodical sequence, maximizing the effectiveness of the recovery and reconstitution efforts and minimizing system outage time due to errors and omissions.

AWS maintains a ubiquitous security control environment across all Regions. Each data center is built to physical, environmental, and security standards in an active-active configuration, employing an n+1 redundancy model to provide system availability in the event of component failure. Components (N) have at least one independent backup component (+1), so the backup component is active in the operation even if all other components are fully functional. To reduce single points of failure, this model is applied throughout AWS, including network and data center implementation. All data centers are online and serving traffic; no data center is “cold.” In case of failure, there is sufficient capacity to enable traffic to be load balanced to the remaining sites.

AWS services alignment with the CSF

In an increasingly interconnected world, applying strong cybersecurity risk management practices for each interconnected system to protect the confidentiality, integrity and availability of data is a necessity.

AWS public and private sector customers fully expect that AWS employs best-in-class security to safeguard its cloud services, and the data processed and stored in those systems. To effectively protect data and systems at enterprise scale, security cannot be an afterthought, but rather an integral part of AWS systems lifecycle management. This

means that security starts at Phase 0 (systems inception) and is continuously delivered as an inherent part of the AWS service delivery model.

AWS exercises a rigorous, risk-based approach to the security of our services and the safeguarding of customer data. It enforces its own internal security assurance process for our services, which evaluates the effectiveness of the managerial, technical, and operational controls necessary for protecting against current and emerging security threats impacting the resiliency of our services.

Enterprise-scale commercial cloud service providers such as AWS are already subject to robust security requirements in the form of sector-specific, national, and international security certifications (for example, FedRAMP, ISO 27001, PCI DSS, SOC, and so on) that sufficiently address the risk concerns identified by public and private sector customers worldwide.

AWS adopts the security high bar across all of its services based on its “high watermark” approach for its customers. This means that AWS takes the highest classification level of data traversing and stored in its cloud services, and apply those same levels of protection to all of its services and for all of its customers. These services are then queued for certification against the highest compliance bar, which translates to customers benefiting from elevated levels of protection for customer data processed and stored in the AWS Cloud.

Conclusion

Public and private sector entities acknowledge the security value in adopting the NIST CSF into their environments. U.S. federal agencies, in particular, are directed to align their cybersecurity risk management and reporting practices to the CSF. As U.S. state and local governments, non-U.S. governments, critical infrastructure operators, and commercial organizations assess their own alignment with the CSF, they need the right tools and solutions to achieve a secure and compliant system and organizational risk posture.

Throughout this whitepaper, we have identified the key features of AWS service offerings that organizations, agencies, and users can use to align to the CSF Functions, Categories, and Subcategories.

This means you can strengthen your cybersecurity posture by leveraging AWS as part of your technology to build automated, innovative, and secure solutions to achieve the security outcomes in the CSF.

Contributors

Contributors to this document include:

- Min Hyun, Sr. Manager, Security/Compliance/Privacy
- Michael South, former Principal Industry Specialist, ADFS-DC Tech
- James Mueller, Sr. Technical Program Manager, Federal Compliance
- Francesco Grande, Solutions Architect
- Carmela Gambardella, Sr. Solutions Architect
- Luca Iannario, Solutions Architect Manager
- Giuseppe Russo, Security Assurance Manager
- Tim Blair, Manager, Commercial Compliance
- Jonathan Randle, Security Business Development

Document revisions

Date	Description
January, 2025	Updated and aligned to the NIST CSF 2.0.
October 12, 2021	Updated.
January 2019	First publication

Notes

1 Cybersecurity standards usage for control systems in organizations worldwide in 2021, <https://www.sans.org/white-papers/SANS-2021-Survey-OTICS-Cybersecurity/>

2 <https://www.gartner.com/en/articles/3-planning-assumptions-for-securing-cyber-physical-systems-of-critical-infrastructure>

3 HIPAA includes provisions to protect the security and privacy of protected health information (PHI). PHI includes a very wide set of personally identifiable health and health-related data, including insurance and billing information, diagnosis data, clinical

care data, and lab results such as images and test results. The HIPAA rules apply to covered entities, which include hospitals, medical services providers, employer-sponsored health plans, research facilities, and insurance companies that deal directly with patients and patient data. The HIPAA requirement to protect PHI also extends to business associates.

4 PHI includes a very wide set of personally identifiable health and health-related data, including insurance and billing information, diagnosis data, clinical care data, and lab results such as images and test results.

5 <https://www.fsscc.org/About-FSSCC>

6 The AWS Well-Architected Framework documents architectural best practices for designing and operating reliable, secure, efficient, and cost-effective systems in the cloud. It provides a set of foundational questions that allow you to understand if a specific architecture aligns well with cloud best practices.

<https://docs.aws.amazon.com/wellarchitected/latest/framework/welcome.html>

7 <https://www.cisecurity.org/controls/>

8 Federal Risk and Authorization Management Program (FedRAMP) is the U.S. government's standardized, federal-wide program for the security authorization of cloud services. FedRAMP's "do once, use many times" approach was designed to offer significant benefits, such as increasing consistency and reliability in the evaluation of security controls, reducing costs for service providers and agency customers, and streamlining duplicative authorization assessments across agencies acquiring the same service.

9 ISO 27001/27002 is a widely-adopted global security standard that sets out requirements and best practices for a systematic approach to managing company and customer information that's based on periodic risk assessments appropriate to ever-changing threat scenarios. ISO 27018 is a code of practice that focuses on protection of personal data in the cloud. It is based on ISO information security standard 27002 and provides implementation guidance on ISO 27002 controls applicable to public cloud Personally Identifiable Information (PII). It also provides a set of additional controls and associated guidance intended to address public cloud PII protection requirements not addressed by the existing ISO 27002 control set.

10 Available online and classroom training can be found at <https://aws.amazon.com/training>. There are also several books covering many aspects of AWS, which can be found at <https://www.amazon.com> by searching for "AWS". AWS whitepapers can be found at <https://aws.amazon.com/whitepapers>.