



Cloud Security Technical Reference Architecture

Coauthored by:

Cybersecurity and Infrastructure Security Agency,
United States Digital Service, and
Federal Risk and Authorization Management Program

August 2021

Version 1.0

Executive Summary

Executive Order 14028, “*Improving the Nation’s Cybersecurity*” marks a renewed commitment and prioritization of federal cybersecurity modernization and strategy. To keep pace with modern technology advancements and evolving threats, the Federal Government continues to migrate to the cloud. In support of these efforts, the Secretary of Homeland Security acting through the Director of the Cybersecurity and Infrastructure Security Agency (CISA), in consultation with the Director of the Office of Management and Budget (OMB) and the Administrator of General Services acting through the Federal Risk Authorization Management Program (FedRAMP), have developed the *Cloud Security Technical Reference Architecture* to illustrate recommended approaches to cloud migration and data protection for agency data collection and reporting that leverages Cloud Security Posture Management (CSPM). This technical reference architecture also informs agencies of the advantages and inherent risks of adopting cloud-based services as agencies move closer to zero trust architecture.

Authority

Executive Order 14028, “*Improving the Nation’s Cybersecurity*” provides at section 3(c) (emphasis added):

As agencies continue to use cloud technology, they shall do so in a coordinated, deliberate way that allows the Federal Government to prevent, detect, assess, and remediate cyber incidents. To facilitate this approach, the migration to cloud technology shall adopt zero trust architecture, as practicable. **The CISA shall modernize its current cybersecurity programs, services, and capabilities to be fully functional with cloud-computing environments with zero trust architecture.** The Secretary of Homeland Security acting through the Director of CISA, in consultation with the Administrator of General Services acting through the FedRAMP within the General Services Administration, **shall develop security principles governing Cloud Service Providers (CSPs) for incorporation into agency modernization efforts.** To facilitate this work:

[...]

Within 90 days of the date of this order, the Secretary of Homeland Security acting through the Director of CISA, in consultation with the Director of OMB and the Administrator of General Services acting through FedRAMP, shall develop and issue, for the Federal Civilian Executive Branch (FCEB), **cloud-security technical reference architecture documentation that illustrates recommended approaches to cloud migration and data protection for agency data collection and reporting.**

Contributing Authors

Cybersecurity and Infrastructure Security Agency

The Department of Homeland Security (DHS) CISA provides support to agencies for evolving and operationalizing cybersecurity programs and capabilities. As the cyber risk advisor for Federal Civilian Executive Branch (FCEB) agencies, CISA seeks to provide enhanced support for agencies adopting cloud services to improve situational awareness and incident response in cloud environments. CISA is responsible for aiding federal agencies, critical infrastructure, and industry partners as they defend against, respond to, and recover from major cyber attacks.

United States Digital Service

The United States Digital Service (USDS), is a senior team of technologists and engineers that support the mission of departments and agencies through technology and design. USDS's multi-disciplinary teams bring best practices and new approaches to support government modernization efforts. USDS is situated under OMB.

OMB produces the president's budget and examines agency programs, policies, and procedures to assess with the president's policies and coordinates inter-agency policy initiatives. OMB evaluates the effectiveness of agency programs, policies, and procedures, assesses competing funding demands among agencies, and sets funding priorities. OMB also ensures that agency reports, rules, testimony, and proposed legislation are consistent with the president's budget and administration policies. OMB also oversees and coordinates the administration's procurement, financial management, information, and regulatory policies. In each of these areas, OMB's role is to help improve administrative management, develop better performance measures and coordinating mechanisms, and reduce unnecessary burdens on the public.

Federal Risk and Authorization Management Program

Established in 2011, FedRAMP provides a cost-effective, risk-based approach for the adoption and use of cloud services by the Federal Government. FedRAMP empowers agencies to use modern cloud technologies, with an emphasis on security and protection of federal information.

FedRAMP is a program under the General Services Administration (GSA), which manages and supports the basic acquisition and procurement functions of federal agencies. GSA supplies products and communications for U.S. government offices, provides transportation and office space to federal employees, and develops government-wide cost-minimizing policies and other management tasks.

Table of Contents

1.	Introduction.....	1
2.	Purpose and Scope	1
2.1	Key Programs and Initiatives	2
3.	Shared Services Layer.....	3
3.1	Cloud Service Models Overview	3
3.2	Introduction to FedRAMP	7
3.3	Security Considerations under FedRAMP	8
4.	Cloud Migration.....	10
4.1	Designing Software for the Cloud	10
4.2	Cloud Migration Strategy.....	11
4.3	Cloud Migration Scenarios	13
4.4	Developing a DevSecOps Mentality.....	17
4.5	Building Scalable, Repeatable Architectures.....	19
5.	Cloud Security Posture Management.....	21
5.1	Defining CSPM.....	22
5.2	CSPM Outcomes.....	24
5.3	Adopting CSPM Capabilities.....	28
	Appendix A – Glossary and Acronyms	40
	Appendix B – Resources.....	42

Table of Tables

Table 1:	Common Cloud Migration Challenges	11
Table 2:	Technical Challenges in Cloud Migration	12
Table 3:	Benefits to Cloud Migration	12
Table 4:	Cloud Migration Strategies.....	13
Table 5:	CSPM Outcomes	30

Table of Figures

Figure 1:	Cloud Security Technical Reference Architecture Composition and Synergies	2
Figure 2:	Responsibilities for Different Service Models	4
Figure 3:	Scenario 1 – Notional Phase 1 Architecture.....	14
Figure 4:	Scenario 1 – Phase 2 Notional Architecture with Out-of-Band Data Transfer	14
Figure 5:	Scenario 2 – Notional Migration of a Website to a PaaS	15
Figure 6:	Scenario 2 – Notional Website with CDN.....	15
Figure 7:	Scenario 2 – Notional Final Architecture of the New Website	16
Figure 8:	Scenario 3 – Notional Deployment of SaaS-based Website Monitoring	16
Figure 9:	DevSecOps Loop.....	17
Figure 10:	Reference Architecture for a Build System with Security Testing.....	18
Figure 11:	Reference Architecture on Centralized Security Services.....	20
Figure 12:	Service Deployments and Integrated Solutions.....	31
Figure 13:	Authentication Realms	33

1. Introduction

Executive Order 14028, *“Improving the Nation’s Cybersecurity”* (May 12, 2021)¹ marks a renewed commitment and prioritization of federal cybersecurity modernization and strategy. Among other policy mandates, the Executive Order embraces zero trust as the desired model for security and tasks the Cybersecurity and Infrastructure Security Agency (CISA) with modernizing its current cybersecurity programs, services, and capabilities to be fully functional with cloud-computing environments. While the Executive Order marks a shift in federal policy, many efforts undertaken in recent years support the key tenets of this Executive Order.

These preexisting efforts should continue; however, new leadership, evolving threats, and changing requirements and technologies present an opportunity to enhance existing strategies and architectural approaches. In addition, recent cyber breaches affecting cloud computing environments have had wide-ranging implications and demand a national response. These compromises demonstrate that “business as usual” approaches are no longer acceptable for defending the nation from cyber threats. Furthermore, cloud migration requires cultural changes, priorities, and design approaches that must be embraced, driven, and supported by the entire organization.

This Cloud Security Technical Reference Architecture builds on the initiatives above and supports the continued evolution of federal agencies within a rapidly evolving environment and technology landscape through a focus on cloud modernization efforts, namely: shared services, designing software in the cloud, and cloud security posture management.

2. Purpose and Scope

The purpose of the Cloud Security Technical Reference Architecture is to guide agencies in a coordinated and deliberate way as they continue to adopt cloud technology. This approach will allow the Federal Government to identify, detect, protect, respond, and recover from cyber incidents, while improving cybersecurity across the .gov enterprise. As outlined in the Executive Order, this document seeks to inform agencies of the advantages and inherent risks of adopting cloud-based services as they move closer to zero trust architecture². The Cloud Security Technical Reference Architecture also illustrates recommended approaches to cloud migration and data protection for agency data collection and reporting.

This technical reference architecture is intended to provide guidance to agencies adopting cloud services in the following ways:

- **Cloud Deployment:** provides guidance for agencies to securely transition to, deploy, integrate, maintain, and operate cloud services.
- **Adaptable Solutions:** provides a flexible and broadly applicable architecture that identifies cloud capabilities and vendor agnostic solutions.
- **Secure Architectures:** supports the establishment of cloud environments and secure infrastructures, platforms and services for agency operations.
- **Agile Development:** supports a dynamic development and engineering cycle that prioritizes the design, development, and delivery of capabilities by building, learning, and iterating solutions as agencies transition and evolve.
- **Zero Trust:** supports agencies as they plan to develop zero trust architectures.

¹ <https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/>

² <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf>

This technical reference architecture is divided into three major sections:

- **Shared Services:** This section covers standardized baselines to evaluate the security of cloud services.
- **Cloud Migration:** This section outlines the strategies and considerations of cloud migration, including explanations of common migration scenarios.
- **Cloud Security Posture Management:** This section defines Cloud Security Posture Management (CSPM) and enumerates related security tools for monitoring, development, integration, risk assessment, and incident response in cloud environments.

While each major section covers unique aspects of cloud security, they share common synergies that support the overall goal of modernizing cloud security. Figure 1 details the composition and commonalities.

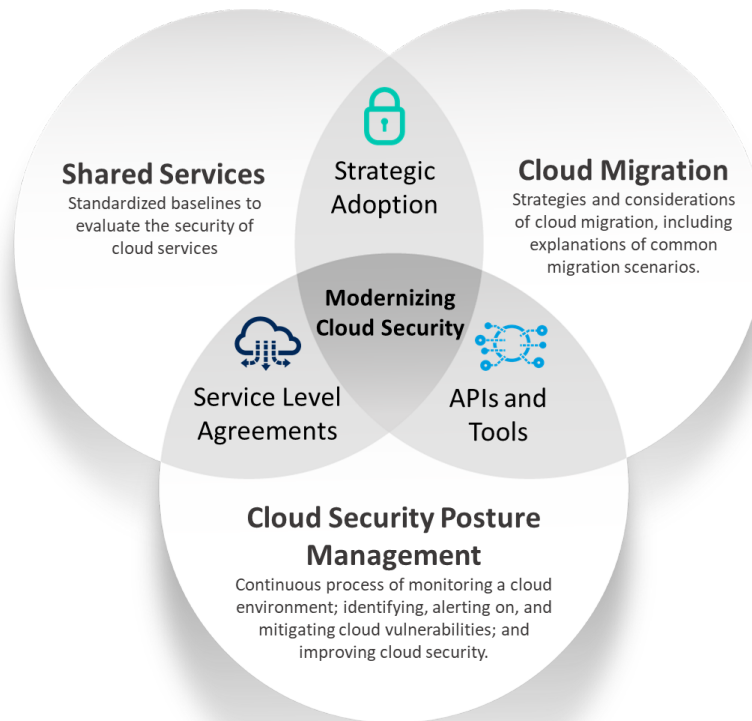


Figure 1: Cloud Security Technical Reference Architecture Composition and Synergies

2.1 Key Programs and Initiatives

The following are key federal cloud programs and strategies in place to ensure both information technology (IT) modernization and cloud security.

Federal Risk and Authorization Management Program

The Federal Risk and Authorization Management Program³ (FedRAMP) was established in 2011 to provide a cost-effective, risk-based approach for the adoption and use of cloud services by the Federal Government. FedRAMP empowers agencies to use modern cloud technologies, with an emphasis on security and protection of federal information.

³ <https://www.fedramp.gov/>

Cloud Smart Initiative

As a successor to the legacy Federal Cloud Computing Strategy “Cloud First”, the Federal Cloud Computing Strategy “Cloud Smart”⁴ was initiated in 2017 as a result of the Report to the President on Federal IT Modernization.⁵ Cloud Smart emphasizes the three pillars of security, procurement, and workforce. While these pillars are still a focus of the cloud strategy, there is a stronger cross-cutting emphasis with security; for example, the emphasis on building expertise in the federal IT workforce should include prioritizing skill sets and training in cloud computing security architectures.

3. Shared Services Layer

This section introduces shared services and the security implications for agencies and vendors. The section provides an overview on cloud services models and explains how agencies can leverage FedRAMP services to support their cloud migration. It is important to note that the features of the cloud services models described in this section rely on contractual terms set during procurement; cloud acquisition is outside of the scope of this technical reference architecture.

This section will:

- **Define cloud service models:** Identify and define cloud service models and how this document uses these definitions in comparison with other authoritative resources.
- **Introduce FedRAMP:** Explain FedRAMP and associated roles and responsibilities.
- **Outline security considerations under FedRAMP:** Describes FedRAMP requirements for continuous monitoring, incident response, and the authorization boundary.

3.1 Cloud Service Models Overview

There are many options when moving infrastructure, applications, or services into the cloud. Typically, these options are referred to as “_aaS” where the “_” can be a letter or a series of letters that describes the type of cloud-based offering. The National Institute of Standards and Technology (NIST) has defined three basic cloud service models: SaaS, or Software-as-a-Service; PaaS, or Platform-as-a-Service; and IaaS, or Infrastructure-as-a-Service.⁶

- **Software-as-a Service (SaaS):** Consumers are users of the provider’s applications running on an underlying cloud infrastructure. Applications are accessible via various client platforms. Consumers do not manage or control the underlying infrastructure.
- **Platform-as-a-Service (PaaS):** Consumers have the capability to deploy custom applications using provider-supplied languages, libraries, services, and tools on the cloud infrastructure. Consumers do not manage or control the underlying infrastructure, but they have control over the deployed applications and potentially the configuration settings of the provider-supplied environment that is hosting the application.
- **Infrastructure-as-a-Service (IaaS):** Consumers have the capability to provision computing resources to deploy and run environments and applications. Cloud providers manage the underlying infrastructure while the consumers have control over the computing resources, including some control of selected networking components (e.g., host- versus network-based firewall).⁷

⁴ <https://cloud.cio.gov/strategy/>

⁵ <https://www.cio.gov/assets/resources/Report-to-the-President-on-IT-Modernization-Final.pdf>

⁶ <https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf>

⁷ <https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf>

As cloud has evolved over the years, there is an ever-growing list of other _aaS acronyms for various offerings including Desktop-as-a-Service (DaaS), Security-as-a-Service (SECaaS), Artificial Intelligence-as-a-Service (AIaaS), Container-as-a-Service (CaaS), Disaster Recovery-as-a-Service (DRaaS), Internet of Things-as-a-Service (IOTaaS), Location-a-a-Service (LaaS), Monitoring-as-a-Service (MaaS), Unified Communications-as-a-Service (UCaaS), and Workspace-as-a-Service (WaaS), among others. These additional offerings overlap with the three basic service models and are blurring the delineation between the three, further complicating responsibilities around maintenance and security.

SaaS, PaaS, and IaaS are the most prevalent cloud services offerings, and each offering has differences in how they are consumed and protected. This is commonly represented via the shared security model, illustrated in Figure 2. Such models outline which party has responsibility for technology, security, data, etc. Figure 2 includes layers for Configuration and Identity and Access Management (IAM) that other reference architectures do not include.

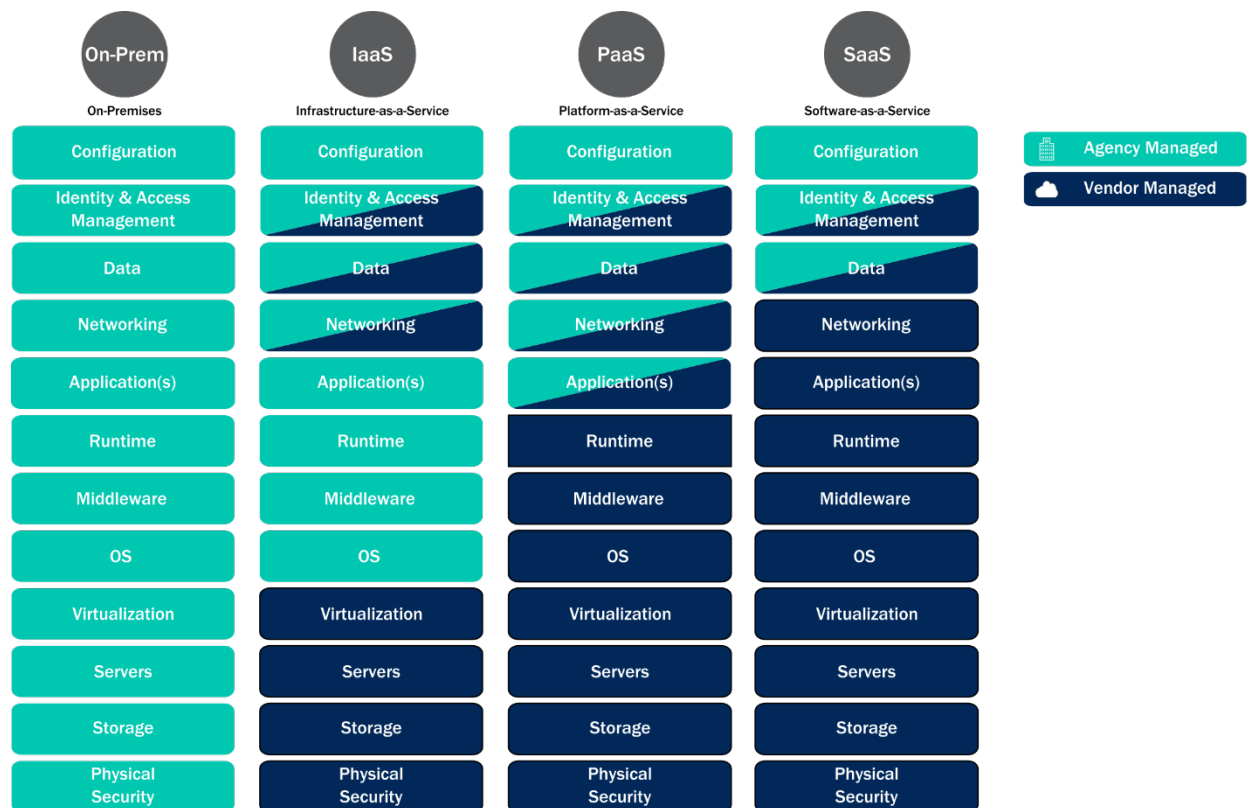


Figure 2: Responsibilities for Different Service Models

The shared services model (Figure 2) shows that the responsibility for securing a SaaS offering relies heavily upon the service provider. However, this also means that the agency consuming the service is placing more trust in the service provider. This contrasts with IaaS, where much responsibility falls on the agency, some responsibility resides with the CSP, and other responsibilities are shared. CSPs may define this shared security relationship differently from one vendor to the next. Agencies must clearly identify and understand the delineation of responsibilities between themselves and their CSP. Agencies should carefully set up service level agreements (SLA) to define expectations and responsibilities with each of their CSPs. Agencies may find that they need to change their security posture to stay current with their CSP(s) as they update service offerings. Agencies should ensure that they properly understand the security posture of their elected CSP(s) both initially and continuously over time.

3.1.1 Cloud Service Options

As mentioned above, there are three primary cloud service options: SaaS, PaaS, and IaaS. Each type of cloud service offers unique features and carries its own security implications that agencies should consider when implementing efficient architectures. Each cloud service is detailed in the subsections below.

3.1.1.1 Software-as-a-Service

SaaS offerings are generally dedicated in nature and target a business need such as communications (e.g., email), document management, or human resources functions. SaaS offerings are typically offered through the web, but they can also be applications or application programming interfaces (APIs) that can be integrated with another service. The hardware and software are controlled by the service provider with few shared responsibilities.

Some SaaS providers will have the ability to integrate with existing identity access providers; others will not have authentication integration options and will have their own identity realm. IaaS and PaaS providers may have some SaaS offerings as part of their portfolio of available services.

3.1.1.2 Platform-as-a-Service

PaaS offerings are often included as part of IaaS but can also be offered independently. In this environment, platforms, such as web servers and databases, are available to build solutions. The advantage of PaaS over IaaS is that agencies can focus on creating services for mission needs rather than buying, deploying, and managing server hardware or the application or database server. This means that an agency can focus on managing platform resources and developing and deploying services and solutions, rather than focusing on the administration of the underlying infrastructure.

3.1.1.3 Infrastructure-as-a-Service

IaaS environments will offer a rich set of services and functions that can be used to build and orchestrate solutions. Agencies should understand and consider features native to the cloud so they can take advantage of these resources when developing solutions. Such features include elasticity and scalability, as well as the virtualization of resources such as networks, operating systems, containers, etc.

3.1.2 Deployment Types

The service offerings described above can be deployed in the cloud in four different ways. The following are the different cloud deployment types and their NIST definitions:

Private: The cloud infrastructure is provisioned for exclusive use of an organization comprised of multiple customers (e.g., an agency with multiple business units). It may be owned, managed, and operated by the organization, an authorized third party, or combinations of them. The infrastructure may exist on-premises with the organization or off-premises with the cloud provider.

Community: The cloud infrastructure is provisioned to a specific community of consumers that have shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more organizations, an authorized third party, or some combination of these entities. The infrastructure may exist on or off premises.

Public: The cloud infrastructure is provisioned for use by the general public. It may be owned, managed, and operated by one or more organizations, an authorized third party, or some combination of these entities. The infrastructure exists off-premises.

Hybrid: The cloud infrastructure is a composition of two or more of the above deployment models (i.e., Private, Community, or Public). In this instance, multiple deployment models are connected through a standardized or proprietary technology offered by the provider to maintain compatibility of data and applications.⁸

Regarding community cloud, many consider government cloud offerings to be a type of community cloud model. While government cloud deployments may offer some protections beyond public cloud offerings, such as US citizens working at the CSP data center, there may be some disadvantages, too. Typically, CSPs offer new security features and tools first to the public model. It may take weeks, months, or years for these same security features and tools to be offered to government cloud deployments. Also, some features within the tools offered by CSPs in a Public cloud deployment may never be implemented in the associated government deployment. Additionally, government cloud deployments are limited to U.S. regions. Some agencies may require a global reach that is best accomplished through a public cloud deployment.

3.1.3 Multi-Cloud

Agencies are likely to operate in a multi-cloud environment where they consume one or more SaaS offerings and one or more PaaS or IaaS offerings. Agencies operating in a multi-cloud environment need to plan for how they can optimize their use of multi-cloud environments while maintaining situational awareness and proper security practices in each CSP they operate within. Agencies can choose to protect each of these services as an entity on its own or they may decide to maintain a holistic view of their security posture for all the services they consume.

Agencies should evaluate how to best monitor each cloud service they use and maintain situational awareness and proper security practices. It is important to find parity in the security information between the different cloud offerings an agency uses. Data normalization of logs by type will help achieve parity as each of the service offerings will have variations in field names and the number of fields in the logs, they make available. Agencies should determine if they will consolidate logs to a central location for analysis and, if so, which logs and how the logs will be backhauled. Some logs will have a consolidated location, like authentication logs if using an integrated identity access provider across multiple CSPs. Additionally, many tools are available within CSPs and by third-party vendors for security analysis within a given CSP or across multiple CSPs. Agencies will want to determine which of these tools best improve their security posture based on their unique needs.

When planning to adopt cloud services agencies must determine how they will implement authentication and access management for each service. They must consider where their identity provider will reside (e.g., on-premises, in a CSP—if they have more than one, which CSP will host the identify provider). Additionally, agencies will need to consider when and how to enable security features such as multi-factor authentication (MFA) or convenience features such as single sign-on.

When operating in a multi-cloud environment, agencies should be cognizant of the potential for vendor lock-in. Vendor lock-in occurs when a tenant has dependencies on services and resources within a CSP. In some cases, choosing to architect solutions that introduce vendor lock-in can provide many advantages. While in other situations, agencies might need to architect solutions with minimal vendor lock-in so that solutions can easily be deployed across different services with minimal changes to configurations and settings for deployment.

⁸ <https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf>

3.2 Introduction to FedRAMP

FedRAMP was established in 2011 to provide a cost-effective, risk-based approach for the adoption and use of cloud services by the Federal Government. FedRAMP empowers agencies to use modern cloud technologies, with an emphasis on security and protection of federal information. FedRAMP is a government-wide program that promotes the adoption of secure cloud services across the Federal Government by providing a standardized approach to security and risk assessment for cloud technologies and federal agencies.

Benefits

- Reduces duplicative efforts, inconsistencies, and cost inefficiencies.
- Establishes a public-private partnership to promote innovation and the advancement of more secure information technologies.
- Enables the Federal Government to accelerate the adoption of cloud computing by creating transparent standards and processes for security authorizations and allowing agencies to leverage security authorizations on a government-wide scale.

Goals

- Grow the use of secure cloud technologies in use by government agencies.
- Enhance the framework by which the government secures and authorizes cloud technologies.
- Build and foster strong partnerships with FedRAMP stakeholders.

3.2.1 FedRAMP's Stakeholders: Roles & Responsibilities

Four stakeholder groups serve roles in FedRAMP—CSPs, Third Party Assessments Organizations, federal agencies, and the Joint Acquisition Board (JAB).

Cloud Service Providers

The Federal Government is one of the largest buyers of cloud technology, and CSPs offer agencies innovative products that help them save time and resources while meeting their critical mission needs. CSPs who have a Cloud Service Offering (CSO) that is being used by the Federal Government should obtain a FedRAMP Authorization. FedRAMP provides a standardized security framework for all cloud products and services that is recognized by all Federal Civilian Executive Branch (FCEB) agencies. CSPs only need to go through the FedRAMP Authorization process once for each CSO and perform continuous monitoring of each authorized service. All agencies review the same continuous monitoring deliverables to create efficiency across the government. The FedRAMP Program Management Office (PMO) provides training, guidance, and advisory support to CSPs, helping them navigate the FedRAMP process and understand the requirements.

Third Party Assessment Organizations

Third Party Assessment Organizations (3PAOs) play a critical role in the authorization process by assessing the security of a CSO. As independent third parties, they perform initial and periodic assessments of cloud systems based on federal security requirements. The Federal Government uses 3PAO assessments as the basis for making informed, risk-based authorization decisions for the use of cloud products and services. During FedRAMP assessments, 3PAOs produce a Readiness Assessment Report (RAR), which is required for the Joint Authorization Board (JAB) Authorization process. While an RAR is optional for agency authorizations, it is highly recommended. For both JAB and agency authorizations, 3PAOs produce a Security Assessment Plan (SAP) and Security Assessment Report (SAR). The SAP and SAR must be submitted to a government Authorizing Official (AO) for authorization.

Federal Agencies

FedRAMP helps federal agencies use cloud services to securely modernize their technology and support their mission. To do this, agencies use FedRAMP’s standardized baselines to evaluate the security of cloud services. Agencies work with CSPs to review the security posture and authorize the CSO for any cloud services that they wish to use. Agencies can review and reuse CSO security packages once they are designated as “Authorized” within the FedRAMP Marketplace by issuing their own authorization to use the product. FedRAMP’s “do once, use many” principle enables agencies to expand the marketplace of secure cloud services available to the Federal Government.

Joint Authorization Board

The JAB is the primary governance and decision-making body for FedRAMP. The JAB consists of the Chief Information Officers from the Department of Defense (DoD), the Department of Homeland Security (DHS), and the General Services Administration (GSA). The JAB is responsible for:

- Defining and regularly updating the FedRAMP security authorization requirements
- Approving accreditation criteria for Third Party Assessment Organizations (3PAOs)
- Reviewing authorization packages for cloud services based on the priority queue
- Granting provisional authorizations for cloud services that can be used as an initial approval that Executive departments and agencies leverage in granting security authorizations and an accompanying Authority to Operate (ATO) for use
- Ensuring that provisional authorizations are reviewed and updated regularly and notify Executive departments and agencies of any changes to provisional authorizations including removal of such authorizations
- Establishing and publishing priority queue requirements for authorization package reviews

The *JAB Charter* provides additional details on the objectives and responsibilities of the board.⁹

3.3 Security Considerations under FedRAMP

FedRAMP’s role is to provide a standardized approach to security and risk assessment for cloud technologies and federal agencies. Even after authorization, CSPs and agencies should be aware of ongoing security requirements and considerations.

3.3.1 Continuous Monitoring

It is inevitable that the security posture of an agency’s system will change after receiving authorization. This may be due to changes in the hardware or software on the cloud service offering or the discovery of new exploits. Ongoing assessment and authorization provide federal agencies using cloud services a method of detecting changes to the security posture of a system for the purpose of making risk-based decisions. The *FedRAMP Continuous Monitoring Strategy Guide* describes the FedRAMP strategy for a CSP to use once it has received a FedRAMP Authorization (via agency authorization or JAB provisional authorization).¹⁰ The CSP must continuously monitor the cloud service offering to detect changes in the security posture of the system to enable well-informed risk-based decision making. The guide instructs the CSP on the FedRAMP strategy to continuously monitor their systems.

3.3.2 Incident Handling

The Federal Information Security Modernization Act of 2014 (FISMA), codified in relevant part at 44 U.S.C. § 3551, et seq., is the authoritative source for incident definitions. 44 U.S.C. § 3552(b)(2) defines

⁹ https://www.fedramp.gov/assets/resources/documents/FedRAMP_Joint_Authorization_Board_Charter.pdf

¹⁰ https://www.fedramp.gov/assets/resources/documents/CSP_Continuous_Monitoring_Strategy_Guide.pdf

an "incident" as "an occurrence that (A) actually or imminently jeopardizes, without lawful authority, the integrity, confidentiality, or availability of information or an information system; or (B) constitutes a violation or imminent threat of violation of law, security policies, security procedures, or acceptable use policies." The terms "security incident" and "information security incident" are used interchangeably with "incident" in this document.

After a CSP obtains a FedRAMP Agency ATO or Provisional-ATO (P-ATO) for its service offering, it enters the continuous monitoring (ConMon) phase. Clear and timely incident communication to relevant stakeholders is a key aspect of ConMon to ensure that all incident handling is transparent, and so that all stakeholders are aware of the current status and remediation efforts. The *FedRAMP Incident Communications Procedures*¹¹ document outlines the steps for FedRAMP stakeholders to use when reporting information concerning information security incidents, including response to published Emergency Directives. FedRAMP requires CSPs to report any incident (suspected or confirmed) that results in the actual or potential loss of confidentiality, integrity, or availability of the cloud service or the data/metadata that it stores, processes, or transmits. Reporting real and suspected incidents allows agencies and other affected customers to take steps to protect important data, to maintain a normal level of efficiency, and to ensure a full resolution is achieved in a timely manner

3.3.3 Authorization Boundary

FedRAMP provides guidance to CSPs for developing the "authorization boundary" associated with their CSO to support their FedRAMP Authorization package.

Authorization Boundary: An authorization boundary provides a diagrammatic illustration of a CSP's internal services, components, and other devices along with connections to external services and systems. An authorization boundary encompasses all technologies, external and internal services, and leveraged systems and accounts for all federal information, data, and metadata that a CSO is responsible for. The authorization boundary is a critical component associated with the NIST Special Publication (SP) 800-37, Guide for Applying the Risk Management Framework (RMF) to Federal Information Systems and OMB circular A-130, Managing Information as a Strategic Resource.

FedRAMP is currently updating the Authorization Boundary Guidance document¹² to reflect changes to cloud computing technology and federal information security policy relevant to FedRAMP. The major changes will include:

- Scoping and defining the Authorization Boundary in the cloud;
- Defining data types, including federal data and federal metadata in the cloud; and
- Leveraging interconnections, external and corporate services.

FedRAMP does provide Continental United States (CONUS) requirements for the data centers, but only for the high baseline. For FedRAMP low and moderate baselines, agencies should be aware that there are no implicit or explicit protections for federal agencies that ensures their data will stay only within the US or that their resources will only be established in regions that operate within the US. Agencies must establish these boundaries and expectations with their CSPs and address any Outside the Continental United States (OCONUS) concerns through SLAs or memorandums of understanding (MOUs).

¹¹ https://www.fedramp.gov/assets/resources/documents/CSP_Incident_Communications_Procedures.pdf

¹² The FedRAMP Authorization Boundary Guidance is open for public comment through September 13, 2021. <https://www.fedramp.gov/blog/2021-07-14-Public-Comment-Boundary-Guidance/>

4. Cloud Migration

This section introduces the compute plane and considerations for agencies as they design, implement, and maintain digital services in the cloud. To ensure an efficient and secure transition to cloud services, agencies should:

- **Design software for the cloud:** Identify the appropriate services and capabilities to implement from the start to create a secure and efficient cloud environment.
- **Create a cloud migration strategy:** Design an agency-specific plan to transition data and services from an on-premises environment to a cloud environment.
- **Adopt a Development, Security, and Operations (DevSecOps) approach:** Create reliable automated digital services by utilizing code and integrating support personnel.
- **Build scalable, repeatable architectures:** Employ best practices and necessary resources to maintain a robust cloud environment into the future.

4.1 Designing Software for the Cloud

Agencies can utilize the flexibility of the cloud to combine services in support of their mission. Agencies should work to implement security measures into their cloud-based digital services as early as possible in the Software Development Life Cycle (SDLC). Agencies that facilitate DevSecOps with automated security testing will develop architectures that are scalable, repeatable, and reliable. This process requires collaboration across agency teams to build the best digital services. DevSecOps can combine with centralized SaaS supported by IT departments to enable well tested software for release.

Cloud-based digital services can span IaaS, PaaS, and SaaS. These service models, along with the on-premises model, vary in who is responsible for different layers of the system architecture, as discussed in Section 3. It is imperative for agencies to confirm the services and functions their vendors are providing and are not providing.

4.1.1 Why Shift Software to the Cloud

Agencies moving software and digital services from an on-premise data center to the cloud can produce more reliable, scalable, and predictable software. Cloud services allow agencies to have disaster recovery available in other geographical areas and quickly expand capacity when needed, all without having to purchase another data center. However, this does not mean existing monolithic services need to be migrated as a starting project. Instead, agencies can start with smaller, internal projects and tools to gain experience and confidence working in this new environment. It is also an opportunity to redesign older digital services to enable bold progress or modernization.

The cloud offers a long list of well-known benefits; in particular, one that agencies should consider is that building zero-trust, and more secure applications, can be an easier lift in the cloud. By looking for the appropriate FedRAMP approval level for services in the cloud, agencies can typically expedite an ATO easing the migration process. Correctly configuring these services, establishing effective IAM roles, and protecting sensitive information using encryption provided by a Key Management System (KMS) may be the responsibility of DevSecOps teams or other administrators. Section 5 has additional guidance for Cloud Security Posture Management.

Agencies should consider the security advantages of using APIs (see Section 5.3.8) or data services to securely manage their cloud deployments. Services from CSPs and third party vendors can provide access to the same data without forcing agencies to build, verify, and maintain complex software. APIs provided by CSPs and others typically have a full staff of developers and other experts who focus solely on these

systems. Creating an equivalent team within an agency can be costly and time consuming, drawing resources away from an agency's mission.

4.2 Cloud Migration Strategy

Cloud migration is the process of moving business operations and missions into the cloud. For many agencies, this means shifting from legacy infrastructure that may no longer support their needs to a modern infrastructure that enjoys the support of a more flexible and more cost-effective solution for an agency's application. Cloud environments inherently involve a shift in mindset from on-premises solutions. Certain cloud functions can operate in ways that on-premises functions cannot, such as Infrastructure as Code (IaC) concepts. These concepts include dynamic provisioning and decommissioning of resources based on the elasticity of demand on services or temporal-based maintenance to replace portions of infrastructure for security purposes.

Cloud migration involves a lot of preparation and depends on the size of the application ecosystem, the age of the current applications and systems, the user base, and the amount of data. Agencies should consider the age and quantity of data in their application ecosystem; data accumulated over time is like a force of gravity, increasing inertia the bigger it gets. When agencies decide to migrate their application ecosystem to the cloud, they should weigh benefits and challenges to adopting cloud-based technologies.

4.2.1 Possible Cloud Migration Challenges

All large-scale software projects have their challenges, but the shift from on-premises to the cloud has some unique aspects around personnel, funding, and data. Table 1 lists common challenges that agencies face when migrating to the cloud.

Table 1: Common Cloud Migration Challenges

Common Challenges	How does it affect the migration?
Funding	The application infrastructure and data may exist in multiple environments for a period of time requiring an overlap in funding needs before cost savings may be realized. Additionally, there are costs associated with transferring data—while moving data into a CSP is often free, it can be costlier to move data out depending on the architecture and approach.
Onboarding	Onboarding should include extra time to train the team on the new technologies used to facilitate a successful migration for their application.
Infrastructure Support	A team without cloud migration experience may need help setting up servers, network support, their application, and database in the cloud.
Staffing	As a project grows, a dedicated team may be needed to focus on supporting the migration effort.
Policy Support	As cloud migration generally pushes the boundary of the existing application/project's ATO, policy support will need to support the project by documenting critical changes to extend the boundary of cloud usage to their application.

In addition to common challenges, agencies should consider technical challenges of data migration. Large amounts of data take longer to migrate, validate, and support. Migration difficulties further increase if there are additional requirements that cause little to no downtime for applications or when the underlying data changes frequently. Table 2 details technical challenges related to migrating data to the cloud.

Table 2: Technical Challenges in Cloud Migration

Technical Challenges	How does it affect the mitigation?
Data Integrity	The migration must ensure the security of the data during the transfer as well as the integrity of the data once it has reached its final location of storage.
Minimizing Downtime	Many applications within agencies are operational during government business hours, allowing a weekend exercise of downtime. Selective applications may have more stringent downtime requirements. When replacing a system, minimizing downtime in the transition requires preparation and, in many recommended cases, an iterative rollout of the application in the cloud.
Network Support ¹³	When a large amount of data passes through an agency's network infrastructure in support of a data migration, the agency should understand latency and throughput aspects of the network. These measurements can drive better decisions on how to migrate the data to the cloud vendor's environment.

4.2.2 Benefits of Cloud Migration

Cloud services offer agencies a range of operational and financial advantages since many business and mission processes are cloud-centric in nature. NIST presents the five essential characteristics of cloud computing in SP 800-145¹⁴ as on-demand self-service, broad network access, resource pooling, rapid elasticity, and measured service. Hardware can be provisioned according to tenants' needs, which represents a fundamental shift away from traditional hardware procurement and management. In addition, tenants may forego instantiating servers altogether (both virtual or bare metal) and build on platforms offered by the CSP. This allows agencies to transfer responsibilities to the CSP, including security and health monitoring, as well as vulnerability and patch management. Provisioned resources may also reside across multiple geographic locations and availability zones within regions, rather than within a single location such as an on-premises server room or data center. When researching different cloud services, agencies should consider their own assets and needs to determine whether cloud services would be appropriate to implement. Table 3 lists notable benefits of cloud migration.

Table 3: Benefits to Cloud Migration

Benefits	How does it benefit a project?
Broader Support	Agencies may choose from a wide range of cloud vendors and support.
Flexibility in Design	Cloud services provide managed services such as document storage, database storage with replication, and application interfaces for automation.
Scalable Performance	Cloud services support a broad range of horizontal scalability, the ability to add more machines to an application's pool of resources.
Cost	CSP services can increase efficiency while allowing agencies to direct financial resources towards mission-critical tasks.
Disaster Recovery	Agencies with off-premises cloud data and infrastructure are better positioned to recover from adverse events at agency offices (e.g., natural disasters).

¹³ Transferring data over an agency's network is only one option. There are other services that can be used to migrate data into the cloud, like copying data to disks and transporting them to the CSP by ground or air.

¹⁴ <https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf>

4.2.3 Cloud Migration Strategies

Table 4 notes some of the major cloud migration strategies popularized by industry partners. Agencies may need to use multiple strategies when migrating an application.

Table 4: Cloud Migration Strategies

Cloud Migration Strategy	Details
Rehost	This technique recreates the application architecture in a “lift and shift” model, shifting the original setup onto servers in the cloud.
Refactor / Rearchitect	This method restructures the application into use cases with the rationale that it will be able to leverage cloud native services from a code and architecture perspective.
Revise / Re-platform	Revising an application will migrate and augment part of an application to utilize cloud native services. A popular solution is to take advantage of cloud native managed databases due to its lower effort to maintain.
Rebuild	Rebuilding an application requires discarding the existing application, and recreating the application utilizing the cloud infrastructure. This relies on creating or situating the application into a cloud native solution.
Replace	This technique eliminates the need of the legacy application by migrating the use cases to a SaaS environment with a third-party vendor.

When migrating to the cloud, agencies may have to account for the nuances of migrating different types of services to and between cloud environments. For example, an agency may choose to migrate development processes. In this case, DevSecOps can be used to maintain newly integrated cloud-native solutions over time and to meet the unique scalability and flexibility needs of on-demand infrastructure. For instance, an agency may decide to leverage containerization to facilitate the orchestration of computing resources for consumers of each service.

4.3 Cloud Migration Scenarios

Agencies encounter unique and complex challenges during migration to the cloud. The following subsections outline common migration scenarios for agencies.

4.3.1 Scenario 1 – PDF Storage Claims to the Cloud (IaaS)

Scenario 1 Description:

An agency is migrating an internal application with 10,000 users where millions of portable document format (PDF) files are uploaded and stored, summing 1 Petabyte of data (1,000 Terabytes). The application uses an on-premises datacenter where the data are stored across multiple server racks.

In Phase 1 of this cloud migration, the agency wants to begin storing new uploaded files in the cloud but has not transferred all the older files. In this scenario, the agency will need an additional layer to manage the identification of stored files locations. The agency should research how to properly redirect newly uploaded files to the cloud environment and should redirect users via a reverse proxy to the proper file location, since files may now be split between on-premises and cloud. Finally, the agency will also need to carefully test all assumptions in a development environment to prepare for the migration.

Figure 3 presents an overview of the architecture for Phase 1.

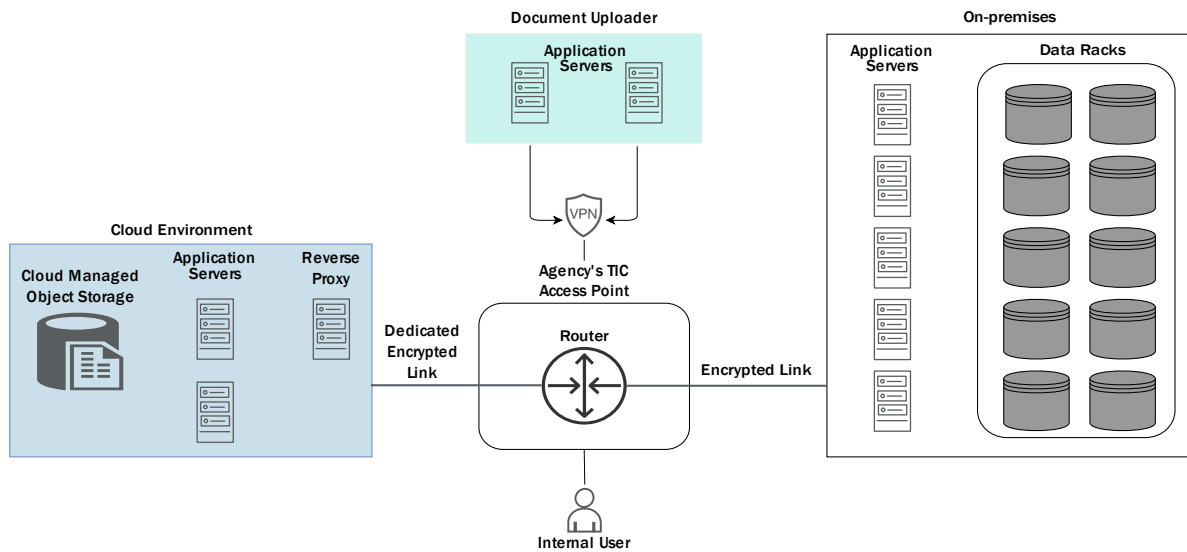


Figure 3: Scenario 1 – Notional Phase 1 Architecture

In Phase 2 of this cloud migration, the agency wants to move the older files to cloud storage. They will need to coordinate with the network team an optimal time to transfer the 1 petabyte of data across the network. Application servers within the on-premises environment will collect the distributed data, generate a set of integrity checksums for future validation, and forward the traffic over encrypted links to the cloud environment. If possible, the agency may consider transferring all data to the CSP via hard drives. This technique may be more efficient than transferring all the data over the network. Figure 4 shows these adjustments.

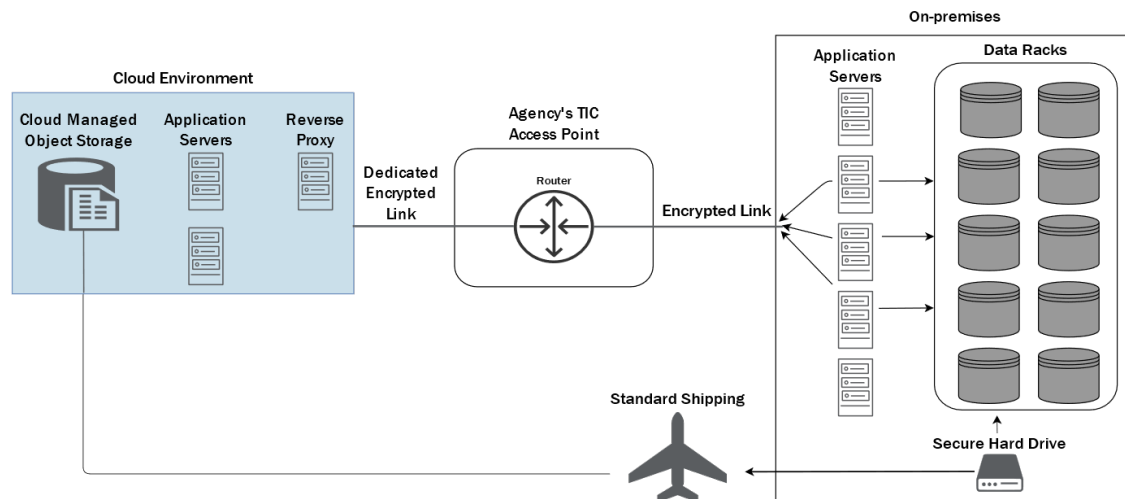


Figure 4: Scenario 1 – Phase 2 Notional Architecture with Out-of-Band Data Transfer

As the data enters cloud storage, it is validated to ensure correctness. Once the data are migrated, the agency should ensure both users and file uploaders are able to seamlessly use the cloud environment. At this point, the on-premises data center can be decommissioned or repurposed.

4.3.2 Scenario 2 –Website Moves to a PaaS Service

Scenario 2 Description:

An agency decides to migrate a legacy website infrastructure hosted on-premises to a modern content management system with a new design. For the past 20 years, an agency has hosted thousands of pages on a locally-maintained, legacy content management system (CMS).

In this scenario, the legacy infrastructure is noticeably dated and many of the web pages require redesign. The agency decides to use a PaaS to build the next enhancement of their CMS. Figure 5 shows the architecture of some of the webpages during the migration and redesign.

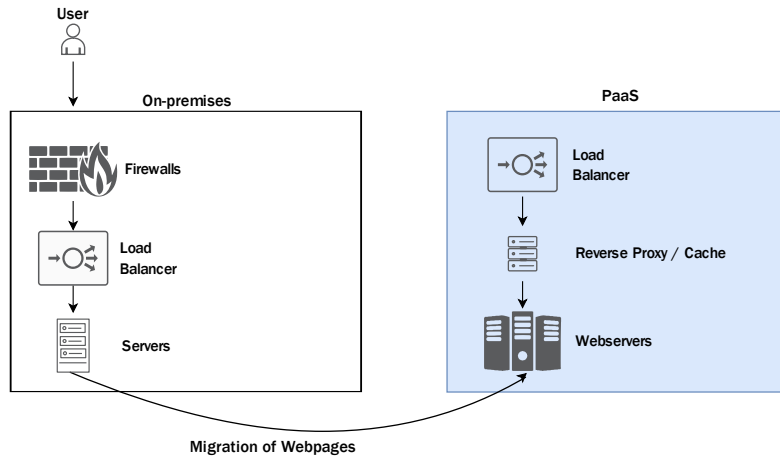


Figure 5: Scenario 2 – Notional Migration of a Website to a PaaS

After the migration and redesign, the agency recognized that most of the webpage content does not change frequently and is suitable for a content delivery network (CDN). Using a CDN will allow the agency to cache most of the content in locations closer to the user, providing faster upload times. The agency will run tests and perform iterative transition files to CDN and configure it to serve user traffic. Figure 6 shows one example of migrating a website to PaaS.

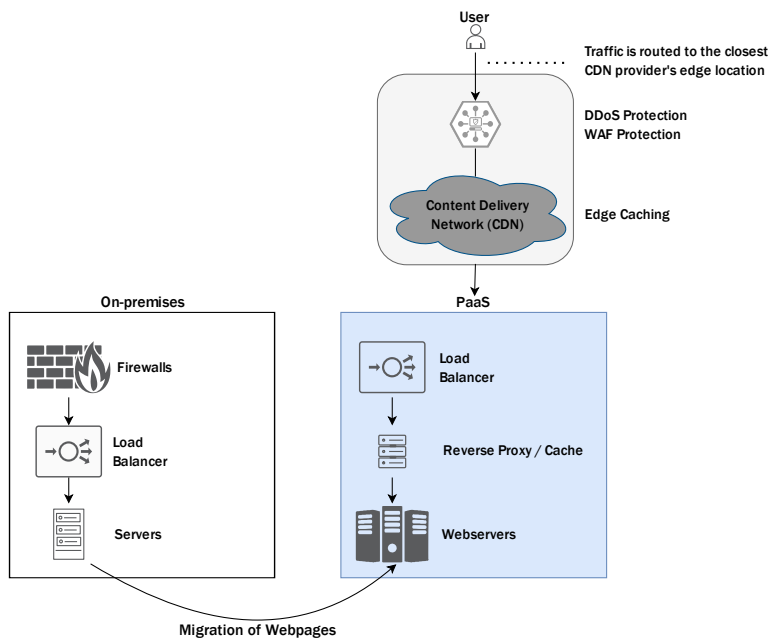


Figure 6: Scenario 2 – Notional Website with CDN

The desired end result will have the on-premises environment decommissioned, and the agency website will be run on the PaaS environment with the CDN entry point as is shown in Figure 7.

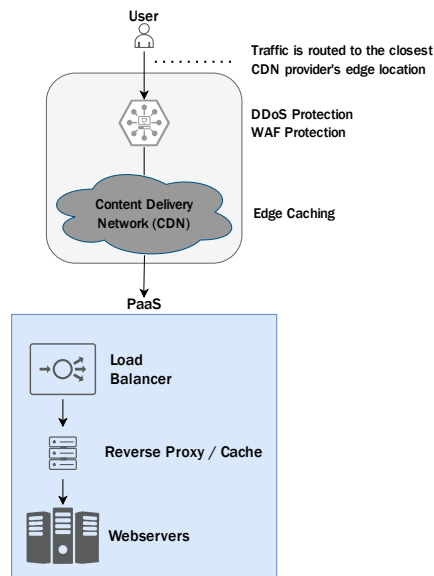


Figure 7: Scenario 2 – Notional Final Architecture of the New Website

4.3.3 Scenario 3 – Monitoring Services for Public Facing Applications

Scenario 3 Description:

An agency is required to monitor its public facing websites for uptime to ensure that it is constantly delivering services for its users.

The agency has multiple websites that are hosted in different locations, so they will need to research monitoring options that can handle the geographically distributed systems. The agency decides on synthetic monitoring, which involves automating potential user actions to see how the system responds and to collect metrics around uptime based on those requests. The agency researches technical considerations and cost tradeoffs of deploying their own monitoring infrastructure in a PaaS or IaaS system versus a SaaS system to collect the metrics. The team settles on using a SaaS system (Figure 8).

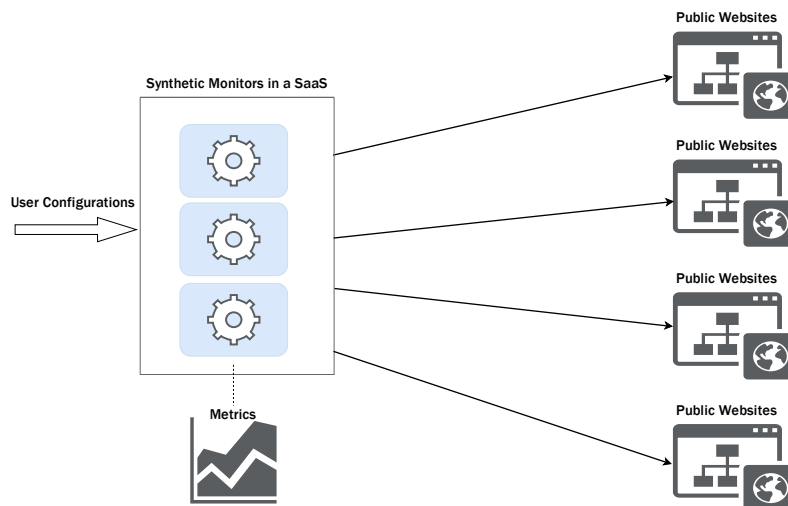


Figure 8: Scenario 3 – Notional Deployment of SaaS-based Website Monitoring

4.4 Developing a DevSecOps Mentality

DevSecOps—a combination of Development, Security, and Operations—is a software development philosophy that tightly integrates writing code with testing, securing, and deploying that code. The traditional DevSecOps loop is illustrated in Figure 9. It can break down silos between the traditional roles of developers, security engineers, operation engineers, and quality assurance professionals and have them function as a team. This is achieved by composing cross-functional teams with these roles working side by side with full ownership for the successful development, launch, and maintenance of their service. DevSecOps should be the primary approach agencies use to develop, secure, and deliver applications in the cloud. DevSecOps often utilizes Continuous Integration (CI), Continuous Delivery (CD), IaC, security testing, and the principle of least privilege to harness automation and produce reliable and predictable digital services that scale.

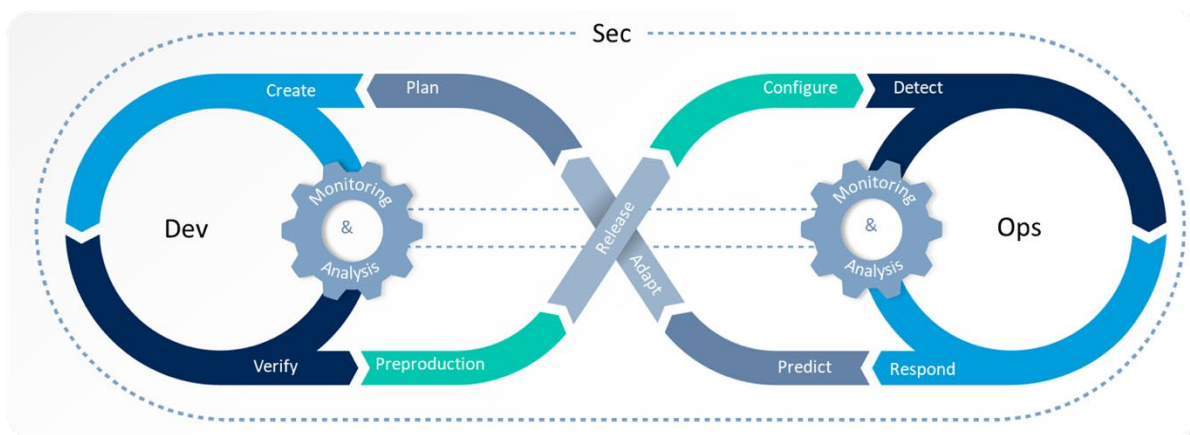


Figure 9: DevSecOps Loop

4.4.1 Continuous Integration and Continuous Delivery

With CI, the repeated activities of code integration, building, and testing are automated to reduce human errors and make the process quick and reliable. This tooling happens earlier in the product lifecycle and is expanded as the project matures. The exact tools used to store source code, build it, and test it vary based on what development teams choose, and there are many options out there including some SaaS products that are FedRAMP approved. IaaS and PaaS providers may also provide these as part of their service.

CD is the process of delivering the code that was integrated, built, and tested on regular intervals using automation. It builds on the CI pipelines to determine when the code is ready for production. Together, these processes are referred to as CI/CD.

By setting the pass and fail criteria for testing, being ready for deployment early in the SDLC, and then using automated processes to check that the criteria are met, agencies end up with more reliable software when it is time for deployment. Not only does this practice catch deployment issues early, but it also increases stakeholder support by supporting an agile workflow that allows for smaller and more frequent course corrections since the partially functioning product can be demonstrated to stakeholders.

4.4.2 Infrastructure as Code

In addition to writing their applications in code, development teams can write their infrastructure as machine-readable definition files that run automated and documented provisioning, runtime changes, and

decommissioning of their digital services. This is known as IaC and it enables teams to review changes to the resources used in IaaS or PaaS before checking in that code. IaC can offer a multitude of benefits:

- Removing the need for a User Interface (UI) which further reduces opportunities for human error,
- Automating deployment of IAM policies, and role based access controls (RBAC), and
- Facilitating security patch deployments and updates.

As with other software, IaC can also perform degrading changes to an environment and possibly introduce new unintended vulnerabilities to a previously secure environment. To reduce the risk of exposure, agencies should monitor IaC code for misconfigurations, and/or perform occasional security code audits for production deployments.

4.4.3 Automated Security Testing

Another factor that can be added to the DevSecOps pipeline is application security testing. This testing as part of the DevSecOps pipeline is a key factor in shifting security "left of boom"¹⁵. Application security testing leverages a combination of static analysis of code that looks for common coding issues like potential Structured Query Language (SQL) Injection vulnerabilities and dynamic testing to see how the code works together. This testing allows agencies to fix potential security issues before they are released into production. In addition to safeguarding the data of the application and mitigating data breaches, application security testing can also reduce the amount of money spent on agencies' Bug Bounty programs the agency may operate (See CISA Binding Operational Directive 20-01¹⁶ for information on vulnerability disclosure programs and bug bounties).

Figure 10 shows a potential architecture for a CI/CD system with security testing in two places. Developers would check their code, both for the application and for the infrastructure, into the appropriate repository. The build system will build the application, and testing begins. Any failed tests would get logged to the monitoring system, and the results will be shared with the developer, possibly with an alert or with a status page. Once all the issues with the build are resolved, the application can be deployed into a development environment for further testing. After all issues are resolved, the application can be promoted to production and is ready to use.

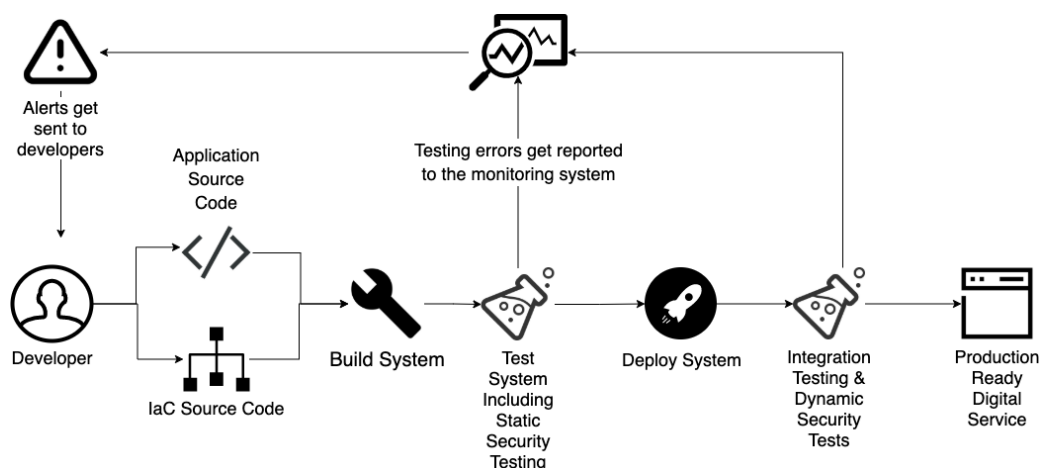


Figure 10: Reference Architecture for a Build System with Security Testing

¹⁵ The idiom “left of boom” in cybersecurity refers to proactive and preventative measures taken to avoid or mitigate a bad event or “boom” (e.g., cyber breach).

¹⁶ <https://cyber.dhs.gov/bod/20-01/>

The testing of infrastructure can also be automated. With definitions of the structure in IaC, security scanners will know what ports are supposed to be open and what are not to identify potential issues early and ideally before they are available on the internet.

4.4.4 Principle of Least Privilege

Agencies should ensure that each DevSecOps team members have sufficient privileges to do their jobs but no more than what is necessary. The Principle of Least Privilege right-sizes the scope and duration of access for each person to perform the duties of their tasks and roles. Some IaaS and PaaS CSPs can facilitate different permissions within the infrastructure based on the activities being done during a given timeframe. When someone oversees operations (a.k.a. On Call), they can be granted additional roles that enable them to access and alter production. Those roles can then be removed when that shift is over. An alternative is a “break glass” procedure to grant temporary access to fix something that is broken.

The risk of ever-expanding roles can also be mitigated with other security best practices, like setting more granular access permissions across the team, and enforcing regular revocations of unneeded access. Procedures for removing access when an employee leaves the team are critical.

4.5 Building Scalable, Repeatable Architectures

Agencies build scalable, repeatable architectures through planning and collaboration. They need to plan out processes, like DevSecOps, that allow teams to iterate quickly, and they need to plan to have the right staff to implement those processes. When agencies collaborate on best practices and cloud functionality like IAM and security monitoring, it allows the development teams to focus on providing the best digital services for their agency and the American people.

4.5.1 Invest in People

Investing in the right people to deliver cloud-based projects is key to a successful project. The three parts to this are training, hiring, and procurement. Federal employees who have been working in traditional software development environments can be re-trained in cloud technology, but this requires agencies to invest in their employees through external classes, trainings, certifications, and the use of work time to study new technology. This can also include training in modern project management methods. To reinforce the trainings, agencies need to allow their employees opportunities to practice their new skills (e.g., through access to sandbox environments that allow for experimentation with these new technologies). Agencies can also leverage programs, like "Interim Ability to Test", that can be used to try a new CSP for a project without having to finish an entire ATO first. Experimentation, iteration, and permission to "fail fast" will help employees that are new to cloud technologies build their skills and deliver superior digital services.

Hiring new federal employees who are already experienced in cloud-based projects is another place to invest in people. It can be challenging to hire for a position that is new because the talent pool is typically limited. One option for hiring technical candidates more efficiently is the Subject Matter Expert Qualification Assessments (SME-QA) process from the Office of Personnel Management (OPM)¹⁷. This allows agencies who need similar staff such as designers or product managers to share job requisitions, filter through candidates using technical assessments, and create a pool of qualified candidates that they can each choose from. SME-QA has increased the number of jobs filled through competitive hiring and reduced the time spent doing it. It can be hard to attract experienced professionals in software

¹⁷ <https://www.performance.gov/cx/blog/CX-hiring-pilot/>

development from the private sector. This is due to a variety of factors, but salary is among them. Agencies can work with OPM to find ways to pay more with the General Schedule (GS) scale, and to offer signing bonuses and quality training opportunities.

Lastly, contractors can be procured to develop digital services and deploy CSP products. TechFAR Hub¹⁸ is a resource for procurement professionals to learn about ways to facilitate the procurement of IT services, including cloud services and contractors who can develop software for it. TechFAR Hub has an initiative called Digital IT Acquisition Professional Training (DITAP)¹⁹ to help procurement professionals learn more.

4.5.2 Document and Share

There are many aspects of developing digital services in the cloud that stay the same from project to project, such as obtaining a domain name for a new website, configuring domain name system (DNS) entries for services, and acquiring Public Key Infrastructure (PKI) keys for transport layer security (TLS). An agency's Chief Information Officer (CIO), or other organizations within an agency, can streamline the work needed to accomplish different tasks by creating, documenting, and publicizing processes. Staff are more likely to follow the procedures when they can find them.

Teams that have been through cloud-based projects should also share their stories, both best practices they identified as well as areas that did not go as well, so that other teams can learn from their experience. Members of these teams can mentor newer teams as well to share the knowledge and skills gained, increasing the overall investment in people.

4.5.3 Centralize Security Services

Agencies should deploy a centrally integrated security services to the greatest extent possible across the enterprise. Fewer separate instances of the same service reduces an agency's attack surface. It also allows teams to reuse these services, save money and time, and specialize on the more interesting parts of their projects. IAM through single sign-on is an ideal place to start, as the CIO likely already has the capacity to enable employees to log in to services, such as email. Even on-premises, Lightweight Directory Access Protocol (LDAP) can broker access to cloud services, reducing the need for employees to remember yet another password. Figure 11 shows a possible configuration with centralized identity and logging.

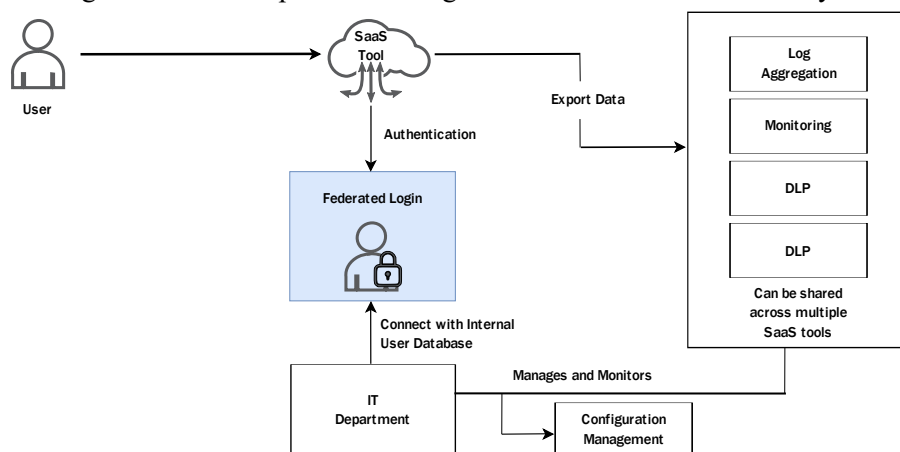


Figure 11: Reference Architecture on Centralized Security Services

¹⁸ <https://techfarhub.cio.gov/>

¹⁹ <https://techfarhub.cio.gov/initiatives/ditap/>

Considerations for LDAP

Recognizing that most agencies rely on on-premises LDAP services, such as Active Directory to access cloud services and resources, agencies are encouraged to work with CSPs to ensure that federated identity services are secured with appropriate logging enabled. CSP LDAP services are evolving rapidly and agencies should continue to work with their vendors and cloud providers to transition LDAP services to the cloud as the primary identity provider.

However, it is possible for centralizing access to CSPs to hinder to other tasks. For example, consider a pilot of a new CSP requires integration into a central IAM system. If a small number of people will be participating in the pilot, and the team responsible for connecting a new CSP to the central IAM system is understaffed, this can delay the use of the pilot. This delay could be avoided by allowing the pilot to use an MFA-based login that may be sufficiently secure for the pilot scope and be done more quickly. While this does not scale to agency-wide products, it can be good for pilots.

Agencies can consider centralizing access to IaaS that is already being used. Agencies can procure the cloud infrastructure in bulk and then give access to different teams as needed, enabling the access desired and helping newer teams to begin using the infrastructure quickly. An agency's cloud team can act more like a PaaS by offering Virtual Machines (VMs) with gold images that match the agency's standard IT configurations. These VMs can be used to install centralized logging and monitoring, and to give access to SaaS products for DevSecOps and application security testing. Together these principles will accelerate the development of digital services in the cloud and save resources.

4.5.4 Grant Staff Appropriate Access

Through onboarding and other documented procedures, agencies should strive to efficiently grant people the access to the CSPs they need to do their jobs and to build good digital services. This is also an important criteria in zero trust architecture as well. When providing teams with the role-based access needed, there is an increased ability to monitor and audit said access. Providing timely access and onboarding also reduces the likelihood that employees will rely instead on "shadow IT" services will be developed that circumvent oversight by IT or security teams, weakening the Federal Government's overall cybersecurity posture.

5. Cloud Security Posture Management

This section introduces CSPM and the related security capabilities and outcomes. This section also highlights some key considerations when migrating to cloud and addresses organizational needs for configuring cloud services and mitigating cloud risks. Additionally, CSPM is contextualized in how such capabilities can facilitate the implementation of zero trust architectures that the Executive Order directs agencies to adopt.

This section will:

- **Define CSPM:** Identify definitions and how this document uses the term in comparison with other authoritative resources.
- **Outline Implementation Needs:** Highlight organizational needs and considerations related to implementing CSPM and zero trust.
- **Harmonize Executive Order Goals:** Provide understanding on the ways in which CSPM supports zero trust goals.

5.1 Defining CSPM

Many networking and cybersecurity terms are commonly used in the context of cloud adoption and operations. Some of these terms have standardized or agreed-upon meanings and definitions. However, many of these terms have divergent definitions and take on different meanings to different stakeholders (e.g., within a given organization, across the Federal Government, within industry, etc.).

The term “Cloud Security Posture Management” has developed relatively recently and is defined differently by various entities. Many of these definitions are similar but written distinctly enough from one another to leave some ambiguity as to the term’s true meaning. Such distinctions of this term’s definition and others may require additional clarification among stakeholders to ensure consensus on their meaning.

For the purposes of this document, CSPM will be used in a broad manner. Namely, CSPM is a continuous process of monitoring a cloud environment; identifying, alerting on, and mitigating cloud vulnerabilities; and improving cloud security. This definition includes the various outcomes (see Section 5.2) and capabilities that support the outcomes (see Section 5.3) identified below.

For the purposes of this document, CSPM capabilities include:

- Security and Risk Assessments,
- Continuous Monitoring and Alerting,
- IAM,
- DevSecOps Integration, and
- Artificial Intelligence (AI)- and Machine Learning (ML)-Based Security Capabilities.

These capabilities seek to support the following activity outcomes:

- Governance and Compliance,
- Policies and Standards,
- Privilege and Identity Access Management,
- Data Protections,
- Infrastructure and Application Protections,
- System Health and Resource Monitoring, and
- Incident Response and Recovery.

Additionally, while this document emphasizes the relationship between cloud adoption and zero trust migration, this does not imply that migrating to cloud services immediately translates into a zero trust architecture. Cloud services *enable* zero trust due in part to the fact that the distributed nature of cloud necessitates additional configuration and management support in order to achieve the kind of security and visibility over assets, users, and data that a zero trust architecture would require.

5.1.1 Why is CSPM Needed?

CSPM provides agencies with access to and management of cloud resources, applications, and data. Agencies moving data and applications to the cloud lose physical access to these deployed resources and must change how they manage governance and compliance requirements for their applications and data. As these cloud deployments mature, they are becoming increasingly more complex, often involving multiple vendors and tools. In addition, recent cyber breaches have had wide-ranging implications; these breaches make clear that proactive management and monitoring are required for defending the Federal Government from cyber threats. Agencies must manage their risk tolerance by continually monitoring and

improving their overall cybersecurity capabilities in a fast-paced environment of evolving threats and where CSPs are constantly changing their product and service offerings.

As agencies migrate to the cloud, there are also opportunities for implementing granular controls and protections, as well as for the management of cloud security by using automated tools for monitoring all aspects of the cloud, discovering threats, and alerting on anomalies. CSPM supports continuous improvement of an agency's cybersecurity posture, and capabilities to enable agencies to keep up with emerging threats, protect against misconfigurations, and reduce the risk of a security incident or data breach.

5.1.2 How can CSPM facilitate Zero Trust?

As directed by the Executive Order, agencies migrating to cloud deployments should apply zero trust principles and transition their environments to zero trust architectures²⁰, commensurate with their risk tolerance. To achieve this, agencies should focus on strengthening fundamental areas of cybersecurity capabilities, such as identity management, asset management, network security, data protection, application security, and visibility, integrated across environments on-premises and in the cloud. CSPM tools and capabilities can be used to support agencies' transition to a zero trust approach to security.

First and foremost, agencies should work towards an identity management solution that provides global identity awareness across cloud and on-premises environments. As agencies migrate services to the cloud, agency users will have identities among a variety of providers. To effectively manage these identities and align security protections holistically, agencies will need to integrate their on-premises identities with those in the cloud environments. Agencies can use CSPM capabilities throughout the identity lifecycle to provide monitoring and analysis and ensure access controls are automatically configured for deployed services.

Agencies should integrate asset and vulnerability management across all agency environments—using automation as much as possible—for cloud deployments. This will require agencies to ensure the integrity of the devices that are used to access services and data, including those in cloud deployments. CSPM tools can be used to gather vulnerability data and to enforce compliance.

In a zero trust architecture, agencies should segment their networks, both on-premises and in cloud environments. Agencies should deploy tools to monitor and provide network visibility into their cloud resources. Agencies can use CSPM capabilities to manage cloud networks and visibility.

Agencies will need to design applications for cloud deployment and consider cloud-native products for application delivery. Agencies should prioritize data and access needs in their design. To this end, agencies should align application security protections based on zero trust principles and integrate their security controls more closely with their application workflows to ensure the protections have the visibility and fidelity needed to provide effective security. Agencies should perform continuous and dynamic application health and security monitoring with external sensors and systems for all applications and services deployed in the cloud. CSPM capabilities can be used for monitoring and managing application deployment configurations.

Lastly, a zero trust architecture demands that agencies adopt a “data-centric” approach to cybersecurity. Agencies should always protect data at rest in the cloud and in transit to and from cloud deployments.

²⁰ <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf>

CSPM tools provide continuous monitoring and alerting on anomalous activity in access logs and help to identify and prevent misconfigurations that may lead to data leakage and data loss.

5.2 CSPM Outcomes

Use of CSPM supports various cybersecurity outcomes, a subset of which are detailed in this section. These outcomes are broadly separated into several categories, corresponding to different security processes that agencies should address. By achieving these different outcomes, agencies can establish strong foundations for the security of their cloud deployments, with protections applied at deployment, during operations, and through post-incident response and recovery.

This section will describe the following outcomes:

- Governance and Compliance,
- Policies and Standards,
- Privilege and Identity Access Management,
- Data Protection,
- Infrastructure and Application Protection,
- System Health and Resource Monitoring, and
- Incident Response and Recovery.

5.2.1 Governance and Compliance

Agencies must comply with both external regulatory and governance requirements and with internally-developed policies and practices. As such, agencies should identify the relevant statutes, regulations, and binding government-wide policies and set in place internal policies and capabilities for assessing compliance. Agencies should ensure compliance extends to all aspects of their cloud services, including acquisition requirements (e.g., obtaining an ATO), billing and contracting renewal, and the termination of services, rather than only deployment and operations. CSPs often natively provide services that comply with many of these regulatory requirements, offering a minimal level of compliance. Many of the solutions also include continual assessment of cloud deployments and environments against these regulatory requirements.

While such services help ensure agency compliance throughout the entire lifecycle of procuring and deploying cloud services, they also require periodic assessment and reassessment. As service providers implement changes and update terms of service, agencies should consider how those changes might natively support compliance with additional regulations or might result in non-compliance and require additional remediation.

5.2.2 Policies and Standards

Beyond governance and compliance, agencies should consider industry standards and best practices to ensure that cloud deployments and services provide a baseline level of operability. Standards and best practices help address the range of deployment requirements, which may differ in areas such as physical security, continuity of operations, and data controls. Again, cloud natively supports many of these outcomes, and agencies should assess which measures satisfy their own requirements and any actions needed to address potential gaps.

This assessment against standards and best practices is not limited to reviewing cloud deployment policies; it should also include policies specific to the cloud service, policies governing non-cloud aspects of an agency enterprise that would intersect with the cloud deployment, and policies for relevant on-

premises services, among others. As service providers implement changes and updates, agencies should continue to reassess and update policies as needed.

This outcome, along with governance and compliance (Section 5.2.1), helps agencies define and set policies to meet their respective requirements. Equally important is how agencies employ these policies and enforce them for cloud services. At the deployment stage, CSPM capabilities can help to ensure these policies are followed in various ways. Approaches, such as IaC or policy-as-code, can enable monitoring, remediation, and automatic enforcement of policies when setting up cloud infrastructure and services. Otherwise, the deployment and enforcement of these policies occur more concretely through other outcomes such as in IAM, data protection, and others.

5.2.3 Privilege and Identity Access Management

One particularly important component of policy enforcement is the handling of permissions and identity. CSPM tools can help integrate privilege and identity access management controls across the entire identity life cycle, as well as provide continuous monitoring and analysis. Monitoring of account activity logs and analysis of behavioral patterns can detect anomalous activity that might indicate a compromise or other potential issues. By consistently managing and defining identity and access controls, CSPM capabilities help ensure that the appropriate privilege and access controls are automatically configured for all deployed services. This addresses weaknesses such as overly permissive access policies and unrestricted code execution privileges, among others.

CSPs are also moving towards natively building in capabilities that support zero trust, which provides a more comprehensive approach to privilege and identity management. This extends the enforcement of privileges and identities to granular account access controls, application and resource authorization, and policy compliance. With these built-in capabilities, cloud then acts as an enabler for agencies in adopting a zero trust approach. However, agencies should still verify that advertised zero trust solutions meet their unique needs for privilege and identity access management.

Agencies should ensure that their own on-premises privilege and identity access controls are up-to-date and consistent with their CSPs' access controls. Best practices such as enabling MFA and setting more granular levels of access and permissions for privileged accounts can limit unauthorized access and privilege escalation within the network, directory services, and applications. Operators should also adopt a posture that enforces least privileges and ensures accounts are kept up to date.

5.2.4 Data Protection

Agencies must work with their CSPs to determine responsibilities for data management and protection. Data protections are required at each stage (create, store, access, roam, share and retire), for all data types (unstructured, structured, semi-structured), and for every state (at rest, in transit, in use) of agency data in the cloud. CSPM capabilities that facilitate policy enforcement can provide various forms of data protections.

Data leakage and data loss are major concerns within data protection. As agencies move data to, from, and within their cloud environments, they must implement and enforce data protection to reduce the potential impact of these risks. Agencies must decide (1) what needs to be protected and at what level, and (2) who controls any data sharing requests and manages the access for each specific CSP. Additionally, agencies should have policies for dealing with data when cloud services are terminated in any way. Access management is also necessary to ensure that deleted data, accounts, and machine images are properly sanitized and inaccessible. An agency's security teams should conduct assessments, enforce controls, and develop analytics for data security and monitoring. This will require knowledge about

ongoing or potential threats, agency policies and management decisions, risk assessments, and vulnerability discoveries.

As described in Section 5.2.1 and 5.2.2, agencies should consider guidance, compliance, policies, laws, regulations, and standards to assess, re-assess and enforce data protection. Agencies should protect intellectual property protection with digital rights management, as appropriate, as well as properly configure cryptographic services such as key management and PKI or symmetric encryption. Re-assessments also should include an examination of the contractual agreement with the CSP. Updates for new and existing features or changes in their service level agreements (SLAs) may be required. For example, in a reassessment, if the CSP already fully encrypts data at rest, assuming compliance and privacy needs are still met, it would be unnecessary for the agency to again encrypt data at rest. However, if the service offering changes, then the service will need to be reevaluated.

Encryption is a key data protection method. Data at rest is often encrypted to avoid data leaks and to protect data in case other security measures fail. Encrypting data in transit allows for data to be transmitted across networks without unauthorized users gaining access. Additionally, depending on their needs, agencies may elect to use client-side or server-side encryption methods. In client-side encryption the agency creates their own key and does not share it, so the CSP cannot view the data being stored. In server-side encryption, the data are encrypted at its cloud destination. Agencies should follow secure key management practices to ensure encrypted data can only be read by authorized parties (see Section 3.5.10). Additional examples of data protection methods to consider include regular and frequent testing of back-ups, separating resources to avoid inadvertent leaks, managing account access, and monitoring of cloud regions, including unused and unsupported cloud regions.

5.2.5 Infrastructure and Application Protection

Infrastructure and application protection can provide security for many layers of cloud usage. These include providing security for the network, resources, and applications associated with an agency's cloud resources. Agencies should deploy vulnerability management procedures and tools to scan their infrastructure, including VMs, virtual networks, applications, containers²¹, and other services used that can be scanned. CSPs have taken steps to natively integrate workload security and posture management logs into management dashboards that allow agencies to create alerts. Some alerts can be responded to automatically, such as triggers that fire to restore altered configurations to an established security baseline. Additionally, third-party tools can also be used to create dashboards for evaluating and assessing cloud security posture. Security management and risk management dashboards and tools allow CSPs to take measurements and compile reports about the effectiveness of decisions. This helps agencies to improve, maintain, or make new decisions as needed. Many of these services are considered proactive and can improve efficiencies in defending against many common and traditional attack vectors. In addition, these infrastructure protections all support agencies' migration to zero trust architectures, providing visibility and analytics into users, devices, network environments, application workloads, and data with automation and orchestration.

Network Protections

Within network protections, proper configuration supports and ensures that networking permissions, segmentation, firewall, proxy, certificates, etc., are correctly configured. Additionally, agencies will take precautions regarding network access and network security settings (e.g., encrypting connections,

²¹ While short-lived containers may not require scanning during execution, container images should be scanned for vulnerabilities in the pre-deployment phase.

authentication, etc.), in accordance with Section 3,d of Executive Order 14028. This aligns with the zero trust tenet that all communication is secured.

Resource Protections

Resource protection, including CSP service configuration protection, is another key component of infrastructure protection. CSP SLA provisions provide protection by implementing some of the items from sections above, such as data protection, and can make CSPs responsible for securing some portion of the provided resources. Other SLA details may enforce physical access protection and monitoring and could include infrastructure location considerations to meet certain requirements. Providing secure access to resources with automated enforcement of policy is a fundamental security capability in a zero trust architecture.

Application Protections

Application protection involves scanning applications running on platforms or middleware, scanning containers, and scanning applications prior to production release or when containers are uploaded to a container repository. Agencies should assess and, where appropriate, limit the degree to which applications can be accessed by other agency resources and vice versa. This identification and mitigation can protect the applications and enable and facilitate quicker responses to their potential misuse. Additionally, strong application security is a key design principle of zero trust.

Vulnerability Management

The management of vulnerabilities, patches, and versions are tied together. By periodically running scans, agencies can ensure that vulnerabilities are systematically discovered and mitigated. This will ensure that keep systems current and patched to required versions and help identify and remove antiquated software. Depending on the architecture used to manage systems, updates may vary. Some updates may be performed in place, while others will follow a vulnerable resource being replaced by a recently patched resource. These updates can be performed without loss of access. Vulnerability management is critical to securing all resources in a zero trust architecture.

5.2.6 System Health and Resource Monitoring

Beyond managing threats to cloud service deployments from malicious actors and activity, CSP tools similarly provide insight into the general operation of the service to ensure proper utilization and system health. For example, indicators such as high central processing unit (CPU) usage or shortages in memory may not be indicative of malicious actors but could point to improper configuration or non-optimal status of services and systems. These tools monitor for indicative events and trigger notifications to the users or automated actions to remediate the situation. These automated actions help ensure that the service is more robust and that resources are sufficient and accessible. In addition to directly handling resource requirements, such as using load balancers to adjust the number of active instances, this monitoring can include broader indicators of the health of the cloud services such as checking billing and payment status, understanding utilization metrics, and tracking the number of users and their amount of activity. Many provider tools provide curated dashboards for visualizing the most intuitive or immediately important areas of concern, to enable continuous visibility into assets and applications. Monitoring the integrity and security posture of all cloud deployments is a fundamental tenet of a zero trust architecture. In addition, this information should be used to continually improve an agency's security posture.

The range of tools available does create some challenges. Agencies must overcome potential fragmentation or lack of integration across multiple solutions from multiple vendors, particularly from third party vendors, as well as account for multi-cloud deployments which may have different specifications and indicators of system health.

5.2.7 Incident Response and Recovery

Agencies should establish and maintain plans to respond to and recover from cybersecurity incidents. Through their management consoles and CSPM capabilities, CSPs and third parties offer a range of response options, including triggering alerts and automated responses to potential risks. These responses enable rapid remediation and prevent further escalation of critical threats. They also allow for more measured responses by human security operations for less immediate threats. The backend cloud infrastructure similarly supports recovery by deploying new resources in place of compromised ones to ensure continuity of service. Immediate disabling of potentially compromised instances can also allow for uncontaminated forensic analysis during post-incident examination.

Incident response and recovery plans are critical to mitigate threats, ensure continuity of service, and retain artifacts for post-event forensic analysis. These plans should account for native CSPM tools and take advantage of cloud capabilities. This could include steps such as ensuring proper-automated response configuration, streamlining access to archived cloud instances, and coordinating with the CSP's incident response plans. Agencies should recognize and understand the differences and challenges associated with incident response and recovery in the cloud. For example, agencies are unlikely to have any access to the physical hardware that their resources reside on. This also includes preparation to perform digital forensic analysis of compromised cloud resources.²² Additionally, agencies should not assume that their data, applications, and infrastructure are automatically backed up because they are using cloud services. Agencies should pre-position capabilities in advance to facilitate response and recovery and should periodically perform audits and inspections as part of keeping their response plans up to date.

5.3 Adopting CSPM Capabilities

Agencies may have existing on-premises infrastructure, data, and processes that they wish to migrate to one or more clouds. While conceptually straightforward, the means by which an agency migrates are nuanced and can be complicated. Existing systems may not be ideal for cloud environments in general or for re-architected cloud-centric solutions. Agencies will need to determine which options are best for their cloud environments. Capabilities such as monitoring, scanning, reporting, mitigation, and other solutions should be evaluated to ensure a sound security posture. This should include adopting CSPM capabilities to achieve the outcomes identified in the previous section, 5.2.

The following section details the general CSPM capabilities available to agencies and their primary functions. However, there are circumstances unique to each agency that will need to be accounted for as agencies move to CSPs and adopt these capabilities. The shared responsibilities model should be used to address the concerns with integrating capabilities across multiple CSPs, so agencies can maintain situational awareness over the security of their interconnected services. This section also explores the ways in which security tools can be deployed independently or as part of an integrated deployment to support delivery of CSPM capabilities.

²² For example, see NIST's "Cloud Computing Forensic Science Challenges" for examples of forensics challenges associated with cloud environments. NISTIR 80006, August 2020.
<https://csrc.nist.gov/publications/detail/nistir/80006/final>

This includes:

- CSPM Capabilities,
- Independent and Integrated Capabilities,
- Managing CSP Accounts,
- Evolution of the Perimeter,
- Visibility and Sensor Placement,
- Monitoring,
- APIs,
- Telemetry and Logs, and
- Deployment, Automation, and Orchestration.

5.3.1 CSPM Capabilities

CSPs offer CSPM capabilities through native services and third parties. While CSPs may allow for integration of third-party capabilities into their services, a CSP may limit or restrict their services from integrating into external third parties. In addition to the traditional infrastructure- and service-level configurations and traditional intrusion detection and prevention systems (IDS/IPS), CSPs enable agencies to integrate CSPM capabilities at scale. Examples include:

1. **Security and risk assessments:** Security assessment capabilities measure policy performance, posture, and compliance. CSPs offer capabilities like traditional on-premises port, service, and configuration scans; however, CSPs may also have integrated capabilities to improve continuous visibility, automate security, and monitor compliance.
2. **Continuous monitoring and alerting:** To provide insight into system resources and data, CSPs can provide monitoring capabilities that enable agencies to record events and other forensic evidence. Logging services should be designed for continuous diagnostic reporting to maximize visibility. As part of the monitoring services, alerts can be established based upon metrics or anomalous behavior. Additionally, third party security information and event management (SEIM) systems can be used to collect, monitor, and alert based upon the logs provided by a CSP.
3. **IAM capabilities:** CSPs provide the ability to perform important functions to either authenticate or connect to third party authentication brokerages. These functions include managing and rotating keys, credentials, and certificates, as well as creating, configuring and monitoring privilege escalation and access to resources. While these capabilities are offered by most CSPs, agencies should understand the nuances and limitations for each of these services before deploying and not assume they are completely secured or meet all compliance requirements.
4. **DevSecOps integration:** Security integration into each component of the DevSecOps pipeline can automate CI/CD with centralized controls. The pipeline may also be improved with regionally specific deployments to respond to ongoing incidents. This includes
 - a. Remediating misconfigurations from both users and automated deployments,
 - b. Redeploying existing infrastructure in response to incidents,
 - c. Monitoring and redirecting traffic through CDNs to extend data visibility and access controls beyond the traditional network perimeter, and
 - d. Using infrastructure as code to incorporate processes and procedures that minimize environmental drift.
5. **AI and ML-based security capabilities:** CSPs can provide AI and/or ML integration to other security capabilities to automate operations, improve performance, and perform analytics on data streams and data stores.

While CSPs may provide many of the above services, agencies may also look to third-party capabilities to expand, supplement, or replace the native CSP offerings. Alongside dedicated third party offerings found in a CSP's marketplace, agencies may integrate external third-party capabilities via Cloud Access Security

Brokers (CASBs), and SECaaS offerings. CASBs can manage access to secure data with record keeping capabilities that use updated encryption keys and log records to regulate access. CASBs can be combined with other capabilities such as log monitoring and AI behavior monitoring to detect abuse of authorized access, mitigate potential incidents, and secure system configurations and software. SECaaS can provide agencies the ability to outsource a portion of their security to a CSP or third party using automated, maintained security capabilities that simplify agency responsibilities.²³ More specifically, the following outcomes, as described in Section 5.2, may be achieved through the CSP-specific CSPM capabilities and their third party and third party-integrated counterparts.²⁴

Table 5: CSPM Outcomes

CSPM Outcomes	Security and Risk Assessments	Continuous Monitoring and Alerting	Identity and Access Management	DevSecOps Integration	AI and ML-based Security Capabilities
Data Protection	X		X		X
Governance and Compliance	X	X	X		
Incident Response and Recovery	X	X	X	X	X
Infrastructure and Application Protection	X	X	X		X
Policies and Standards	X	X	X	X	
Privilege and Identity Access Management	X	X	X	X	X
System Health and Resource Monitoring	X	X		X	

Agencies should consider how integrated they want to be with each CSP they use. Using integrated services from a CSP can provide benefits for both creating and deploying services and for monitoring and protecting the cloud environment. Native CSP capabilities can benefit from a CSP's own internal testing and improved integration with the same CSP's other capabilities. However, there may be times when tools provided by a CSP do not meet the needs of an agency. In these situations, the agency should evaluate third party tools from either the CSP marketplace or commercial off-the-shelf (COTS) solutions to bridge the gaps. Agencies operating in multi-cloud environments may want to use capabilities that span

²³ While CASBs and SECaaS are typically third-party, CSPs may also offer them natively.

²⁴ Not all CSPs will be able to offer all of these capabilities to achieve the outcomes listed. Agencies should evaluate each of the capabilities offered by a CSP to understand what capabilities are provided and identify gaps that may exist. Furthermore, this mapping may change as CSPs deploy new features for capabilities in the future.

their CSP accounts for a holistic, integrated approach. This can be useful for both deployment and security operations.

5.3.2 Independent and Integrated Capabilities

Agencies' security postures may either be developed around the use of independent capabilities (i.e., non-integrated) and/or integrated capabilities across cloud deployments to better identify existing vulnerabilities and on-going compromises, and to prevent future breaches. Figure 12 depicts separate examples of independent capabilities and integrated capabilities applied to notional deployment pipelines, respectively. This figure notionally represents the shift in control to and from security capabilities with the notion of validated, unvalidated, and unmonitored action flows. Both types of capabilities may help secure each component of a service deployment throughout its lifespan. In addition, these capabilities are able to modify a pipeline component based upon the level of control an agency has over that capability (see Section 3.1 for the Shared Responsibilities Model), so agencies typically maintain the same level control over their deployment pipeline with integrated capabilities but can vary control as needed via independent capabilities.

When independent capabilities are applied in the pipeline, there is little interaction between them. This is shown via the separation between the Vulnerability Scanning and Assessment (VS/A), CASB, and IDS capabilities in the deployment pipeline along the top of Figure 12. This approach gives agencies the freedom to select and deploy capabilities as they see fit. However, in order to gain a holistic view across their deployed capabilities, agencies will need to use a third party tool or deploy their own solutions.

Alternatively, deployments may also be wholly based on integrated capabilities to handle unified coordination across services, such as Security Information and Event Management (SIEMs). This approach natively provides enhanced visibility across the deployed capabilities and across multiple deployment pipelines; however, it may provide less freedom to agencies to deploy the capabilities of their choosing. The notional deployment pipeline along the bottom of Figure 12 displays an integrated set of scanning, authentication, and logging capabilities being applied to different portions of an agency's cloud deployment.

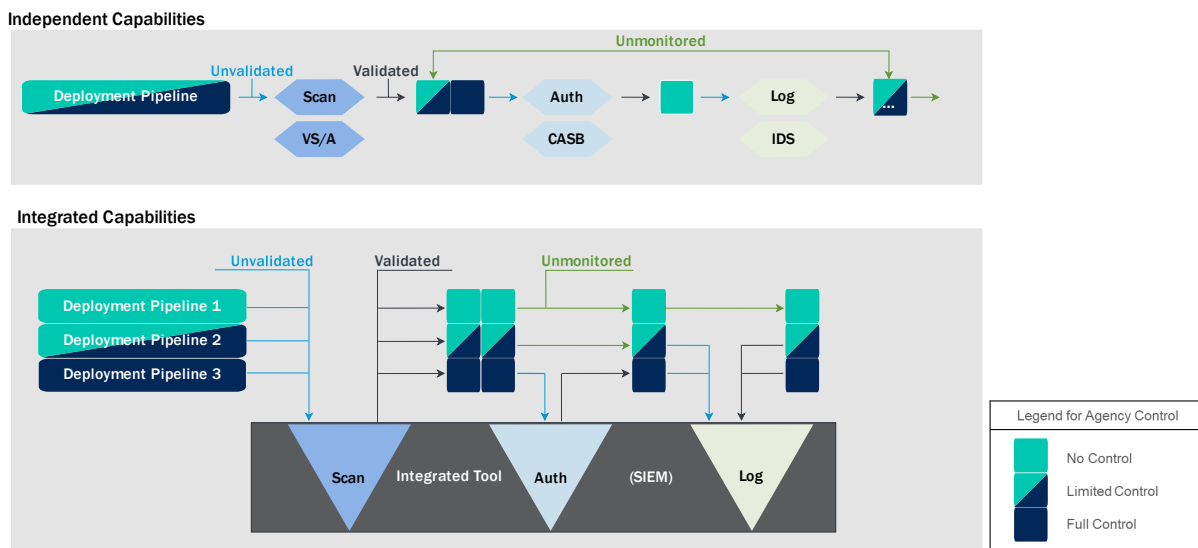


Figure 12: Service Deployments and Integrated Solutions

5.3.3 Managing CSP Accounts

Many CSPs offer the ability to sign up for accounts on demand, while others require coordination on agreements prior to account creation. CSPs differ in whether multiple accounts within a CSP can be linked together, and, if so, *how* these accounts can be linked together. Some CSPs allow for a primary account that can monitor other accounts held by a tenant. Other CSPs use an organization structure within a primary account that allows for unique entities to operate with their own subscriptions, users, and roles. Regardless of how a CSP handles account hierarchy and linkage, the ability to monitor multiple accounts from a single account can provide a holistic view of all accounts and a detailed view of any given account or organization unit.

Agencies may face other considerations regarding account organization and structure. Agencies should be aware that CSPs offering both commercial and government cloud services have not had the ability to move data internally between the two realms. This means that if an agency has a commercial and a government account with a CSP and the agency want all log information for accounts in a single location, they would be required to pay to have the data leave one of the accounts (either commercial or government) so the data could be delivered to the monitoring account. Many CSPs can also create direct network connections to on-premises environments if an agency wants to pay for that capability and bring all security data on-premises for analysis and monitoring.

Agencies should consider creating multiple accounts with a CSP or using built-in account hierarchy tools within a CSP to separate entities within their organization in order to restrict access to assets within a given account. Agencies should then develop criteria to establish the organizational structure for accounts and for granting accesses.

Agencies should also create a plan for how they will establish accounts with CSPs for development and testing. For example, by using IaC, production environments can be replicated quickly so that developers can confidently test code prior to release.

5.3.4 Identity and Access Management

Identity Management

One of the first architecture decisions agencies must make when moving into the cloud is how and where authentication will be performed. CSPs offer both native (e.g., siloed) authentication and integration with identity providers. In many instances federated identity providers are used so that users authenticate to a single identity provider when accessing multiple CSPs, such as email hosted in SaaS and an application hosted by an agency in their IaaS. A federated identity provider can also provide authentication services for users accessing on-premises resources. Some authentication services can integrate MFA and/or single sign-on (SSO). However, while many authentication providers may offer MFA, the MFA may not meet requirements that must be used for government systems, like Personal Identity Verification (PIV)-enabled MFA. In some instances, third party MFA applications can be added to an authentication service, but they will come with additional fees, and some may require the purchase of physical hardware tokens or the use of virtual hardware tokens.

Agencies should carefully manage the different authentication realms that they will use in their environments. An authentication realm is any unique form of authentication that allows a user, process, or system to access another process or system. For example, in Figure 13, there are three resources in a notional IaaS cloud environment: a webserver, a database server, and a fileserver. Each of these are hosted on VMs. The cloud administrator can access these resources through a federated identity provider. This provider can reside in that cloud, in another cloud, or on-premises. The VM server and database

administrators use a username and password to access the server and the database they manage, respectively. The webserver and the server administrator use a certificate to access the web server and VM server they manage respectively. The end users access the web server using a username and password. In this example, there are four distinct authentication realms that are identified by the oval outlines in Figure 13. Because the resources overlap, an exploitation within one authentication realm can lead to malicious activity in resources outside of that authentication realm.

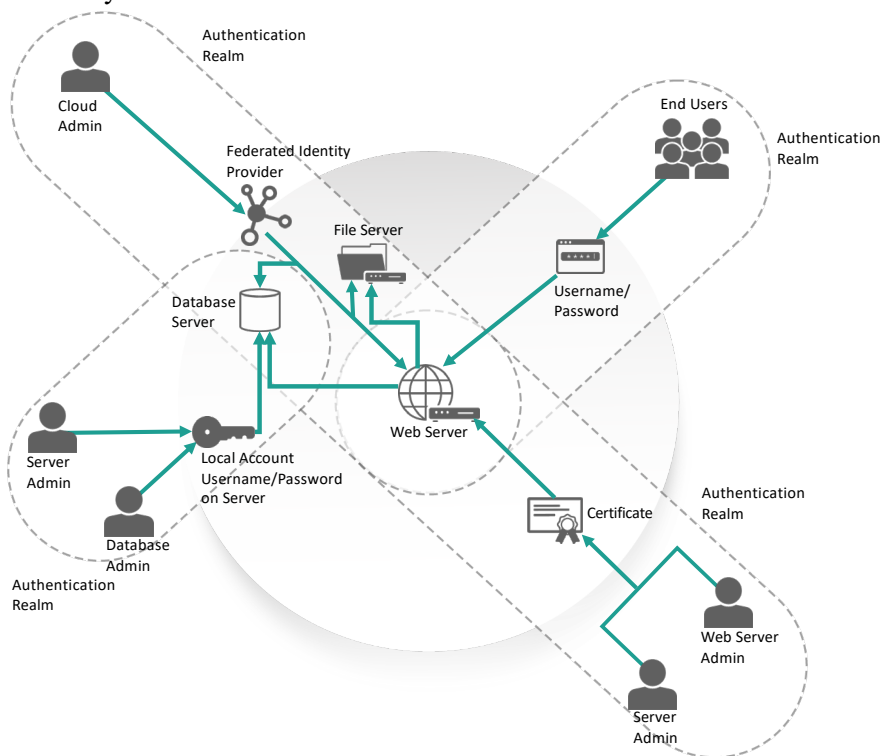


Figure 13: Authentication Realms

Agencies seeking to limit the number of authentication realms, as notionally shown in Figure 13, can use PaaS, which eliminates the underlying server that hosts the database and web server. The web server and database administrators would instead authenticate through the federated identity provider.

Access Management

As agencies adopt a zero trust approach to security, they should enforce least privileges within each authentication realm. This should include distributing access to agency assets (e.g., computer, network, administration, data) and to individual accounts in such a way as to limit the amount of access any one account has to the minimum necessary for that individual or account to perform their responsibilities.

5.3.5 Evolution of the Perimeter

Traditionally on-premises solutions rely on strong perimeter defenses such as firewalls and access control lists. As agencies move into the cloud, their assets cannot be protected by this *castle and moat* paradigm. Agencies will likely operate in a multi-cloud environment where they have varied levels of control over perimeters, which in turn may not be well-defined. In IaaS environments, agencies will be able to emulate traditional network defenses but in a SaaS environment such protections are likely unavailable.

Administrative access to a CSP's console (e.g., via either the web or command line) is open to the internet and agencies do not have the ability to apply allow lists or deny lists to IPs or ports to the CSP console.

However, agencies *can* implement security controls for resources residing within virtual networks created within a CSP.

Instead of taking an outside-in approach to protection of resources, agencies should consider an inside-out approach as part of their perimeter evolution, where protections are put in place focusing on data and then applications and services. Additionally, data, applications, and services can be segmented within the network to enforce more granular security policies for access to the resources.

5.3.6 Visibility and Sensor Placement

When migrating to the cloud, agencies need to understand the limitations of sensor placement and how these limitations may affect their visibility into log data, events, attacks, and other incidents. Sensors can include everything from hardware taps to logs generated from tools like firewalls. Traffic to a tenant on a CSP flows through networking controlled by the CSP. As traffic traverses this path, CSPs conduct their own analyses and may mitigate or eliminate potential threats, such as a Distributed Denial-of-Service (DDoS) attack. These CSP-internal protections can impact an agency's full understanding of the threats facing their cloud resources by limiting tenant visibility. Many CSPs offer a capability to mirror network traffic but only the traffic that makes it to a virtual network inside the CSP can be mirrored. If the traffic is dropped for any reason by the CSP, then an agency will have no visibility of that traffic. Additionally, while CSPs provide protections they generally do not inform the tenant of actions taken on malicious or suspected malicious traffic.

Agencies should consider sensor placement for both north/south traffic (i.e., inbound/outbound) and for east/west traffic (i.e., traffic between services within a network). All north/south traffic should be monitored as it enters and exits the tenant. Furthermore, agencies should consider monitoring east/west traffic especially traffic that moves through peered networks.

5.3.7 Monitoring

Agencies may achieve a robust monitoring capability through the combination of CSPM capabilities, including continuous monitoring and alerting, security and risk assessments, and DevSecOps. Such capabilities can verify compliance, scan vulnerabilities, verify system availability under simulated conditions, identify misconfigurations, and remediate incidents. More specifically, these monitoring capabilities can identify and enumerate service uptime, quality of service, content integrity to detect malicious behavior, and help analysts and engineers investigate compromises and maintain uptime during incidents.

In addition, agencies will need to consider the threat models and geographical deployments appropriate for their operations to effectively monitor cloud resources. If such monitoring capabilities are deployed across multiple geographic regions, agencies may use a unified CSP interface and/or conglomerate multiple regions in a third party service. Once configured, monitoring services should ensure alignment between reported monitoring data and operational cycles like updates and patches. As with integrated capabilities, CSP specific monitoring may promote vendor lock-in, though existing locked-in services may often only be monitored via their respective CSP monitoring services. This may inhibit agencies' abilities to complete any of the previously identified monitoring outcomes (e.g., compliance verification, vulnerability scanning, misconfiguration identification, and incident remediation). Third party monitoring services can provide improved situational awareness across cloud resources, particularly in a multi-cloud case. However, since CSPs may not make all relevant monitoring data available to their users or third parties, these same third party monitoring services may not have the same depth of visibility into an individual cloud environment.

5.3.8 Application Programming Interfaces

A notable departure from on-premises environments is the abundance of APIs in the cloud. While they are not necessarily required, APIs can provide enhanced capacity to use various cloud services and functions. Agencies can also adopt an API-centric and/or microservices approach to their cloud deployments to gain additional benefits (e.g., automation and efficient controls), best practices that minimize environmental drift, and more.

Agencies may increase their attack surface by adopting APIs because they include externally developed code which agencies do not have control nor visibility over. Thus, appropriate security policies should be implemented to mitigate potential cybersecurity risks associated with their usage. Agencies should implement API versioning to keep records of API changes and manage these changes over time. CSPs should also implement a versioning scheme for their APIs, and agencies should verify that such security measures are in place for those API services. Subsequently, by establishing API versioning, CSPs should allow tenants sufficient time to transition between version releases. Agencies can leverage one or more of the following techniques to improve their security posture as it pertains to APIs.

- Use transport encryption to protect confidentiality and integrity of API inputs and outputs.
- Use API access keys as identifiers; these can be used to log which users make certain API calls. To complement this method, agencies should also develop and properly implement an API key revocation policy in the event of API key compromise, along with a corresponding key reissuance policy. Keys should be held in secret, but also be disposable on demand.
- Implement API authorization to enforce user permissions for API calls.

API-centric Architecture

There may be several ways to establish APIs as part of a cloud-based solution. CSPs extensively use APIs and they expose API families as building blocks for tenants for activities such as administration, logging and monitoring, and architecting and interfacing with services. Agencies should evaluate how they can both leverage CSP APIs and build their own API families for the services they implement in the cloud. Great care should be taken to implement adequate security to ensure correct permissions and access to both CSP APIs, and APIs built and implemented by an agency.

When creating applications or APIs that will be made available to customers, agencies should plan how and where they will collect telemetry and how their logs will be made available. Telemetry should be considered for security, performance, errors, connections, etc. Agencies should also include versioning of both the APIs used to collect logs as well as the data structures of the logs they use.

Microservices

Agencies may choose to adopt a microservice architecture approach to their development and production. This is an architectural approach that implements cloud-native applications as a collection of independent, lightweight services. Microservices evolved from the concept of service-oriented architecture (SOA), which in turn arose from monolithic service deployments.²⁵ While the application code for monolithic servers typically takes less development time and is simpler than for SOA or microservices, the tight coupling between its processes provides significantly less robustness when handling application errors and does not scale well. As such, microservices provide an appropriate response to the demands of cloud-based infrastructure.

²⁵ A monolithic architecture is a traditional approach to hosting applications and providing services to an organization. It consists of a single physical server that hosts all processes coupled as a single service.

The deployment of microservices couples well with container technologies due to their lightweight resource consumption and easy deployment, allowing for lower resource usage than when introducing additional full-featured VMs or hardware (e.g., monolithic deployments). An agency can use a container management system to keep track of its microservices with organizational growth. Microservices also enable scalability of applications from a more granular perspective, as each service can be scaled according to its respective load, rather than scaling the entire application around a single bottlenecked service in a monolithic model.

Microservices leverage APIs for inter-service communication. Often, the agency may choose to consolidate these communications into an API gateway; this layer provides a unified interface for an agency to manage security, deployment, analytics, and other service usage. Since microservices are independently deployed and developed, it becomes increasingly useful to leverage an API gateway as more microservices are introduced.

Given that one of the primary benefits of implementing microservices is to reduce overall time and effort in the application development phase, this microservices approach complements with the DevSecOps capabilities of CSPM. Each microservice can be developed separately and follow the CI/CD pipeline of development, deployment, testing, security, and automation. An agency can also implement cloud security monitoring and other functions as microservices and scaled according to operational needs.

Agencies should be aware that adopting a microservices architecture will likely require both a technological and a cultural change in the way that each agency develops software applications. In addition to changes in code structure, the underlying process requires new ways of thinking about software development lifecycles, particularly if shifting services from a monolithic deployment. Agencies also should avoid oversimplification of microservices through excessive functional division, as this overcompensation would create more overhead and undermine the return-on-investment of this paradigm.

5.3.9 Telemetry and Logs

Agencies must understand what logs and telemetry are available to them when consuming cloud services. A systematic review of log management processes is crucial to set up the foundation for monitoring and alerting. Agencies should understand:

- which types of logs are available,
- what data fields are in collected logs,
- when logs are delivered, and
- how collected logs will be processed, stored, and retrieved.

This can help agencies better manage log generation so security teams can more quickly access the logs they need to conduct their operations. Agencies should also take steps to validate and verify that the logs they capture are accurate and are stored appropriately (e.g., in warm storage for on-hand analysis versus cold storage for longer term retention).

When collecting logs from SaaS, PaaS, or IaaS cloud instances, agencies should comply with the logging requirements issued by OMB pursuant to Section 8 of the Executive Order 14028, *“Improving the Nation’s Cybersecurity.”* This provides a list of requirements to improve the ability of federal agencies, CISA, and the Federal Bureau of Investigation (FBI) to hunt for threats and vulnerabilities on federal cloud deployments. To meet this end, agencies can follow some general guidelines.

Time Synchronization

Agencies should ensure all collected logs meet minimum requirements and correlate to the same time zone and the same synced clock. This will allow correlation of all logs from an agency despite regional or

provider differences. Agencies should be aware of and understand the latency of logs collected and made available by the CSP. For example, many CSPs have a latency of up to 15 minutes, which limits real time analysis. Moreover, some telemetry and log collection require action by an agency to receive them, such as installing logging agents on VMs. When collecting logs in multiple regions and time zones, agencies need to understand how each log's time-related fields work. Agencies should verify which time zone each log is captured in, both when in use and when collected. Configurations may be required to default all log timestamps to a specified time zone. If that is not possible, then normalization of log data on ingestion may be performed to ensure accurate querying of events. Additionally, agencies should test for drift in clocks used for creating or reporting time and should engage their CSPs to understand how they ensure accurate timestamps of logs.

Consolidation and Centralization

Agencies should note version numbers associated with collected logs and telemetry, so that if there are new versions, they can perform a comparative analysis of the differences and plan for any necessary changes. Many logs should be configured to automatically be collected and delivered to either storage locations or integrated monitoring capabilities (either CSP provided or third party). Regardless of how collection occurs, and regardless of regional or provider differences, logs should eventually be consolidated in a central location. Some CSPs also allow logs from multiple accounts to be delivered to a primary account which allows for a single location to monitor logs from all accounts. Some of these integration services that cross regions may incur additional costs and agencies should carefully plan for how they will handle logs collected from multiple regions or from multiple CSPs.

Considerations for API Provisioning

When creating applications or APIs that will be made available to customers, agencies should plan how and where they will collect telemetry and how their logs will be made available. Telemetry should be considered for security, performance, errors, connections, etc. Agencies should also include versioning of both the APIs used to collect logs as well as the data structures of the logs they use.

Considerations for SaaS

For SaaS providers, log collection can be performed in several ways. Logs can be made available via an associated IaaS or PaaS account, through API calls to collect logs, by using third party collection tools, and through the export of logs. Exporting logs using a manual process should be avoided, if possible, in favor of an automated scalable collection solution. An alternative option is to keep logs and telemetry in the SaaS environment and use analysis tools offered by the SaaS provider. Because the service provider is responsible for the technology stack and the SaaS offering, tenants do not have the ability to collect additional log data for security purposes other than what the service provider offers. Logs in SaaS environments are typically generated from API calls used by the service provider to build the SaaS offering and they are usually grouped by API families. Access to logs is generally through APIs developed by the service provider, but some service providers may offer security dashboards or log viewers as part of their administrator console. Many SaaS providers build their offerings on top of other offerings from other CSPs. This may limit data that is available to the SaaS provider and therefore limit data availability to the tenant.

Considerations for IaaS

In an IaaS deployment, many logs are available by the CSP that can be captured to gain situational awareness of the environment. These can include network flow logs, API call logs/service event logs, access and identity logs, and health logs. Most IaaS providers have native tools to capture logs and to deposit them into a central location. There may also be options available to collect and share logs across

related accounts so that one account within a CSP can monitor multiple accounts used by an agency. This allows for accounts to be created based upon roles or functions.

5.3.10 Deployment, Automation, and Orchestration

The dynamic nature of the cloud enables agencies to orchestrate services and automate deployment together in ways that cannot be done on-premises. Agencies can automate deployments of new software by incorporating DevSecOps in their development processes. This paradigm fosters a security-first mindset which is especially needed to manage the challenges introduced by CSPs' regular changes to cloud services.

Integrating DevSecOps

DevSecOps is the collaboration of development, security and operation teams encompassed as an integrated unit to achieve the best in developing and deploying code with security built-in from the beginning rather than added on later. While DevSecOps is traditionally geared to production cloud deployments, this security-first mindset is broadly applicable to any cloud environment.

Developers use CI to build and test their deployments. Operation engineers implement CD mechanisms to orchestrate their deployments and monitor them to ensure that they are available and healthy. Security engineers develop tests that run as part of unit, system, and/or integration tests to certify the new deployments meet security standards. In collaboration, security teams work with developer teams during the design process to ensure appropriate security practices are applied, and they also work with operations teams to ensure the deployment is secure, properly monitored, and patched in a timely manner. Throughout the cyclical DevSecOps process, the security team monitors for security issues.

Deployment Management

The virtual environment of the cloud allows agencies to quickly, and fluidly, change components of their cloud deployments. In typical on-premises environments, patches to vulnerabilities and updates to operating systems (OS) and applications happen in-place. Usually, this process results in some down time and is executed outside of standard business hours. Many CSPs and third-party vendors offer tools that change this paradigm, by enabling "zero-downtime" upgrades.

To accomplish this, agencies can create base or "golden" OS images and container images. These images go through processes where required patches and updates are applied, security policies are configured, and security applications are installed. Scans are then executed to verify the results of the process and validate whether an image is secure. Post-creation, these images can then be put into stored repositories and used later to replace running production images. This creation process can be completed on a regular basis so that new images are released monthly, weekly, or even hourly. In addition, system and integration tests should be re-validated so that new updates to applications (OS) or services (container) on golden images do not regress.

An example of this type of deployment might be a container that is built nightly to include the latest libraries that it requires for operation. The container can be run through a battery of tests and security scans then deployed if it passes these tests. All new connections can then be directed to the new container. As existing connections to the previous container terminate, the previous container is decommissioned. If the new container fails a pre-deployment test, then, depending on which test(s) failed, the appropriate engineers are alerted, and they can address the highlighted issue(s).

The cloud also allows agencies to relinquish many maintenance tasks to the CSP, by offering PaaS, SaaS, and Serverless computing options. This can enable agencies to focus on their mission needs. In the container example above, an agency could use a serverless platform offered by a CSP to deploy its

containers. In this case, the agency does not need to worry about various aspects of deployment, such as server acquisition, installation, configuration, OS installation, licensing, patching, monitoring, updating, and the container orchestration software licensing and installation. However, the agency may still be required to perform some configuration of the container software orchestration application.

Key Management

Applying modern cloud-first strategies for key management can enable frictionless encryption across an agency's cloud deployment. Agencies can choose to utilize CSP provided server-side encryption (SSE) or apply a third-party key management service. Agencies are advised against writing their own encryption software. However, before deciding on any key management provider, agencies should ensure the provider meets the requirements of their threat model. Should they find that a CSP or third-party provider does not meet their requirements, agencies should seek to use their own key management strategy.

For example, an agency may want to ensure that the data collected by their application is secured in a way that only the agency can open and view the data and the CSP is unable to access the data. In addition to keys, secrets required for services (e.g., databases, network file shares, APIs, etc.) should be rotated on a periodic basis. Agencies may seek to use offerings by CSPs and third-party vendors that will allow for rotation of passwords, certificates, and keys.

Configuration Management

With rapid deployment available in the cloud, agencies should monitor for unintended configuration changes, i.e., drift, in their environments. A large configuration change is likely to be noticed and detected quickly, but small, incremental changes can easily go unnoticed. Eventually these drifts can compound and create significant changes to an environment such that the environment is no longer compliant with the security plans and ATO for which it was initially approved. Planned changes must be approved to ensure that rogue or unintended changes can be detected and remediated.

Appendix A – Glossary and Acronyms

Application Programming Interface (API): A system access point or library function that has a well-defined syntax and is accessible from application programs or user code to provide well-defined functionality.

Authentication Realm: Any unique form of authentication that allows a user, process, or system to access another process or system.

Authority to Operate (ATO): An official management decision given by a senior organizational official to authorize operation of an information system and to explicitly accept the risk to organizational operations (including mission, functions, image, or reputation), organizational assets, individuals, other organizations, and the Nation based on the implementation of an agreed-upon set of security controls.

Authorization Boundary: All components of an information system to be authorized for operation by an authorizing official and excludes separately authorized systems, to which the information system is connected.

Cloud Access Security Brokers (CASBs): A software tool that manages access to secure data with record keeping capabilities that use updated encryption keys and log records to regulate access.

Cloud Security Posture Management (CSPM): A continuous process of monitoring a cloud environment; identifying, alerting on, and mitigating cloud vulnerabilities; and improving cloud security.

Cloud Service Provider (CSP): An external company that provides a platform, infrastructure, applications, and/or storage services for its clients.

Content Delivery Network (CDN): An interconnected network that pushes caches of files or services across multiple locations to enable secure, fast, efficient delivery of data.

Continuous Integration (CI): The process of automating and integrating modification of code from across multiple teams during software development.

Continuous Delivery (CD): The process of sending new software into production rapidly and automating application delivery.

Continuous Monitoring (ConMon): A process that ensures CSPs continuously maintain the security of their FedRAMP-authorized systems by providing the Joint Authorization Board (JAB) and Authorizing Officials (AOs) monthly insight into the security posture of the system.

Development, Security, and Operations (DevSecOps): A software development philosophy that tightly integrates writing code with testing, securing, and deploying that code.

Digital Services: A generic term to designate applications/services responsible for the delivery of digital information (i.e., data or content) and/or transactional services (e.g., online forms, benefits applications) across a variety of platforms, devices and delivery mechanisms (e.g., websites, mobile applications, and social media). Synonymous with CSP services.

Federal Civilian Executive Branch (FCEB): A subset of U.S. federal departments and agencies that excludes the Department of Defense and agencies in the Intelligence Community.

Identity and Access Management (IAM): A fundamental and critical cybersecurity capability ensures the right people and things have the right access to the right resources at the right time.

Infrastructure-as-a-Service (IaaS): The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run its own software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, and deployed applications; and possibly limited control of select networking components (e.g., host firewalls).

Infrastructure as Code (IaC): The process of managing and provisioning an organization's IT infrastructure using machine-readable configuration files, rather than employing physical hardware configuration or interactive configuration tools.

Intrusion Detection and Prevention Systems (IDS/IPS): Software that automates the process of monitoring the events occurring in a computer system or network and analyzing them for signs of possible incidents and attempting to stop detected possible incidents.

Least Privilege: A design principle whereby each entity is granted the minimum system resources and authorizations that the entity needs to perform its function

Multi-Factor Authentication (MFA): An authentication system that requires more than one distinct authentication factor for successful authentication. Multi-factor authentication can be performed using a multi-factor authenticator or by a combination of authenticators that provide different factors.

Platform-as-a-Service (PaaS): The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services, and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment.

Public Key Infrastructure (PKI): The architecture, organization, techniques, practices, and procedures that collectively support the implementation and operation of a certificate-based public key cryptographic system. Framework established to issue, maintain, and revoke public key certificates.

Service Level Agreement (SLA): A service contract that defines the specific responsibilities of the service provider and sets the customer expectations.

Software-as-a Service (SaaS): The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through either a thin client interface, such as a web browser (e.g., web-based email), or a program interface. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

Telemetry: Artifacts derived from security capabilities that provide visibility into security posture.

Zero Trust: A collection of concepts and ideas designed to minimize uncertainty in enforcing accurate, least privilege per-request access decisions in information systems and services in the face of a network viewed as compromised.

Zero Trust Architecture: An enterprise's cybersecurity plan that utilizes zero trust concepts and encompasses component relationships, workflow planning, and access policies. Therefore, a zero trust enterprise is the network infrastructure (physical and virtual) and operational policies that are in place for an enterprise as a product of a zero trust architecture plan.

Appendix B – Resources

- Cloud Computing PaaS Enterprise Design Pattern § (2010).
https://www.ea.oit.va.gov/EAOIT/docs/April2017docs/041117_EDP_Cloud-Computing-PaaS-EDP-v1.pdf.
- Continuous Diagnostics and Mitigation Program § (2020).
https://www.gsa.gov/cdnstatic/CDM%20Tech_Cap_Vol_Two_Req_Catalog_2020_RFinal_10_2%20.pdf.
- “Digital Services Playbook.” The Digital Services Playbook - from the U.S. Digital Service. Accessed July 9, 2021. <https://playbook.cio.gov/>.
- Department of Defense Enterprise DevSecOps Reference Design § (2019).
https://dodcio.defense.gov/Portals/0/Documents/DoD%20Enterprise%20DevSecOps%20Reference%20Design%20v1.0_Public%20Release.pdf.
- “Enterprise Architecture Quick Guide.” Cloud Security Alliance, 2011.
https://downloads.cloudsecurityalliance.org/initiatives/eawg/EAWG_Whitepaper.pdf.
- Gartner Inc. “How to Protect Your Clouds With CSPM, CWPP, CNAPP and CASB.” Gartner, May 6, 2021. <https://www.gartner.com/en/documents/4001348/how-to-protect-your-clouds-with-cspm-cwpp-cnapp-and-casb>.
- Gartner Inc. “Innovation Insight for Cloud Security Posture Management.” Gartner, January 25, 2019.
<https://www.gartner.com/en/documents/3899373/innovation-insight-for-cloud-security-posture-management>.
- Gwosdz, Medi Madelen. “The Rise of the DevOps Mindset.” Stack Overflow Blog, June 22, 2020.
<https://stackoverflow.blog/2020/06/10/the-rise-of-the-devops-mindset/>.
- Lui, Fang, Jin Tong, Jian Mao, Robert Bohn, John Messina, Lee Badger, and Dawn Leaf, NIST Cloud Computing Reference Architecture § (2011).
<https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication500-292.pdf>.
- Mell, Peter, and Timothy Grance, The NIST Definition of Cloud Computing § (2011).
<https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf>.
- National Cybersecurity Protection System (NCPS) Cloud Interface Reference Architecture § (2020).
https://www.cisa.gov/sites/default/files/publications/CISA_NCPS_Cloud_Interface_RA_Volume-1.pdf.
- “Program Basics.” FedRAMP. Accessed July 2021. <https://www.fedramp.gov/program-basics/>.
- Rose, Scott, Oliver Borchert, Stu Mitchell, and Sean Connelly, Zero Trust Architecture SP 800-207 § (2020). <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf>.
- Young, Lindsay, Aidan Feldman, Mark Headd, Clint Troxel, Waldo Jaquith, Adam Kendall, Britta Gustafson, et al. “18F Blog.” 18F. Accessed July 2021. <https://18f.gsa.gov/tags/devops/>.