



Cloud Security Technical Reference Architecture

Coauthored by:

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

June 2022

Version 2.0

Revision History

The version number will be updated as the document is modified. This document will be updated as needed to reflect modern security practices and technologies.

Table 1: Revision History

Version	Date	Revision Description	Sections/Pages Affected
1.0	August 2021	Initial Release	All
2.0	June 2022	Response to RFC Feedback	All

Executive Summary

Executive Order 14028, “*Improving the Nation’s Cybersecurity*” marks a renewed commitment to 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 implement to zero trust architectures.

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

CISA is the operational lead for federal civilian cybersecurity and executes the broader mission to understand and reduce cybersecurity risk of the nation. In this role, 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	2
2.1 Key Programs and Initiatives	3
3. Shared Services Layer.....	4
3.1 Cloud Service Models Overview	4
3.2 Introduction to FedRAMP	8
3.3 Security Considerations under FedRAMP	11
4. Cloud Migration.....	13
4.1 Designing Software for the Cloud	13
4.2 Cloud Migration Strategy.....	14
4.3 Cloud Migration Scenarios	17
4.4 Developing a DevSecOps Mentality.....	22
4.5 Centralizing Common Cloud Services.....	25
4.6 The Human Element	29
5. Cloud Security Posture Management.....	30
5.1 Defining CSPM.....	31
5.2 CSPM Outcomes.....	33
5.3 Adopting CSPM Capabilities.....	38
6. Conclusion	54
Appendix A – Scenarios	56
Appendix B – Glossary and Acronyms.....	61
Appendix C – Resources.....	64

Table of Tables

Table 1: Revision History	i
Table 2: Common Cloud Migration Challenges	15
Table 3: Technical Challenges in Cloud Migration	15
Table 4: Benefits to Cloud Migration	16
Table 5: Cloud Migration Strategies.....	17
Table 6: CSPM Outcomes	40

Table of Figures

Figure 1: Cloud Security Technical Reference Architecture Composition and Synergies	3
Figure 2: Responsibilities for Different Service Models	5
Figure 3: Scenario 1 – Notional Phase 1 Architecture	18
Figure 4: Scenario 1 – Phase 2 Notional Architecture with Out-of-Band Data Transfer	19
Figure 5: Scenario 2 – Notional Migration of a Website to a PaaS	20
Figure 6: Scenario 2 – Notional Website with CDN.....	20
Figure 7: Scenario 2 – Notional Final Architecture of the New Website	21
Figure 8: Scenario 3 – Notional Deployment of SaaS-based Website Monitoring	22
Figure 9: DevSecOps Loop.....	22
Figure 10: Reference Architecture for a Build System with Security Testing.....	24
Figure 11: Reference Architecture on Centralized Security Services.....	28
Figure 12: Service Deployments and Integrated Solutions.....	42
Figure 13: Authentication Realms	44
Figure 14: PaaS Authentication Example	44

Figure 15:Federated Identity Management 56
Figure 16:Microservices 58
Figure 17: Cloud Warm Site Synchronization and Fail Over Movement..... 59

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, Executive Order 14028 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 Executive Order 14028 marks a shift in federal policy, many efforts undertaken in recent years support the key tenets of this Executive Order. For example:

- Executive Order 13636, *“Improving Critical Infrastructure Cybersecurity”* (February 2013)² expands information sharing programs such as the Enhanced Cybersecurity Services to provide classified and unclassified cyber threat information to U.S. companies.
- Executive Order 13800, *“Strengthening the Cybersecurity of Federal Networks and Critical Infrastructure”* (May 2017)³ authorizes agencies to leverage the NIST CSF to implement risk management measures for mitigating the risk of unauthorized access to government information technology (IT) assets. Executive Order 13800 also directs agencies to prioritize shared services in IT procurements. In this way, Executive Order 13800 prioritizes effective risk management and IT modernization in equal measure, directing agencies to implement effective protections for data while migrating to cloud environments. Executive Order 13800 places increased emphasis on the importance of the CSF and lays the foundation for more rapid cloud adoption across the Federal government.
- Executive Order 13873, *“Securing the Information and Communications Technology and Services Supply Chain”* (May 2019)⁴ emphasizes protections for critical infrastructure IT by securing supply chain acquisition. In this way, it highlights the significance of supply chain and IT procurements for government operations and agency mission fulfillment.

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 in order to succeed.

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

¹ Office of Management and Budget, “Executive Order on Improving the Nation’s Cybersecurity,” (2021), <https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/>.

² Office of Management and Budget, “Executive Order – Improving Critical Infrastructure Cybersecurity,” (2013), <https://obamawhitehouse.archives.gov/the-press-office/2013/02/12/executive-order-improving-critical-infrastructure-cybersecurity>.

³ Office of Management and Budget, “Presidential Executive Order on Strengthening the Cybersecurity of Federal Networks and Critical Infrastructure,” (2017), <https://trumpwhitehouse.archives.gov/presidential-actions/presidential-executive-order-strengthening-cybersecurity-federal-networks-critical-infrastructure/>.

⁴ Office of Management and Budget, “Executive Order on Securing the Information and Communications Technology and Services Supply Chain,” (2019), <https://trumpwhitehouse.archives.gov/presidential-actions/executive-order-securing-information-communications-technology-services-supply-chain/>.

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 Executive Order 14028, this document seeks to inform agencies of the advantages and inherent risks of adopting cloud-based services as they begin to implement zero trust architectures⁵. 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.
- **Development, Security, and Operations (DevSecOps):** supports a secure and 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 adopt zero trust architectures.⁶

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. Understanding the features of shared services and the delineation of responsibilities for managing and securing such services is critical to agencies' cloud migration and security posture management. Migrating to the cloud can help agencies keep pace with the evolving technology landscape by improving both their operations and their security. Lastly, CSPM capabilities will allow agencies to dynamically protect their cloud resources both at scale and across their infrastructure.

Figure 1 details the composition and commonalities.

⁵ National Institute of Standards and Technology, "NIST Special Publication 800-207: Zero Trust Architecture," (2020), <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf>.

⁶ Office of Management and Budget, "Moving the U.S. Government Toward Zero Trust Cybersecurity Principles," (2022), <https://www.whitehouse.gov/wp-content/uploads/2022/01/M-22-09.pdf>.

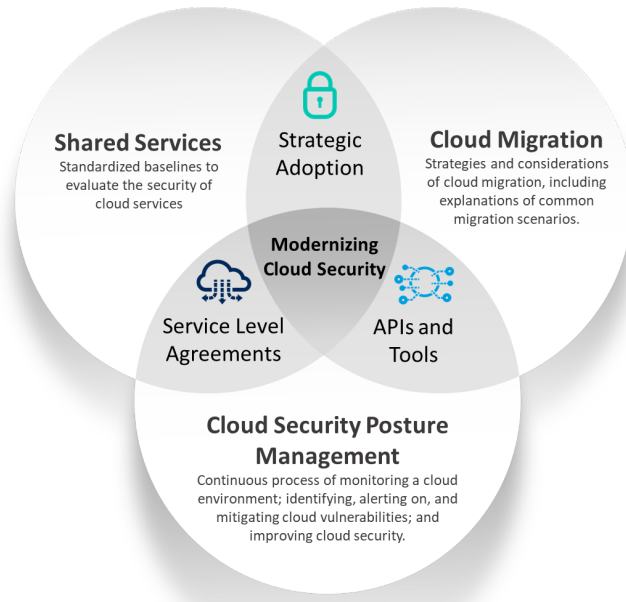


Figure 1: Cloud Security Technical Reference Architecture Composition and Synergies

Appendix A provides three scenarios to highlight considerations associated with the use of federated identity management, microservices, and a warm standby site in the cloud. Appendix B provides a glossary of terms and acronyms found in this technical reference architecture and Appendix C includes a selection of additional resources.

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 Program⁷ (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.

Cloud Smart Initiative

As a successor to the legacy Federal Cloud Computing Strategy “Cloud First”, the Federal Cloud Computing Strategy “Cloud Smart”⁸ was initiated in 2017 as a result of the Report to the President on Federal IT Modernization.⁹ 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

⁷ General Services Administration, “Federal Risk and Authorization Management Program (FedRAMP),” <https://www.fedramp.gov/>.

⁸ Federal CIO Council, “Federal Cloud Computing Strategy: From Cloud First to Cloud Smart,” <https://cloud.cio.gov/strategy/>.

⁹ Federal CIO Council, “Report to the President on Federal IT Modernization,” (2017), <https://www.cio.gov/assets/resources/Report-to-the-President-on-IT-Modernization-Final.pdf>.

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 service 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. 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.¹⁰

- **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).¹¹

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

¹⁰ National Institute of Standards and Technology, “NIST Special Publication 800-145: The NIST Definition of Cloud Computing,” (2011), <https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf>.

¹¹ National Institute of Standards and Technology, “NIST Special Publication 800-145: The NIST Definition of Cloud Computing,” (2011), <https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf>.

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 SaaS, PaaS, and IaaS, further complicating responsibilities around maintenance and security.

However, SaaS, PaaS, and IaaS are the most prevalent cloud service models, and each 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: Responsibilities for Different Service Models

The shared security 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 cloud service provider (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.

Agencies may also use services provided by other agencies, such as a sub-agency using services offered by a parent agency. These services can range from SaaS applications like email to an IaaS environment that the sub-agency is granted access to by the parent agency. In these cases, coordination of roles and responsibilities must be understood between the parent and sub-agency including, but not limited to,

incident management; log monitoring and analysis; identity, access, and credential management (ICAM); and configuration management.

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. Agencies should also be aware that CSPs who offer IaaS services typically also offer PaaS and SaaS services, while CSPs who offer PaaS typically also offer SaaS services. Thus, it is not uncommon for an agency use multiple cloud service models from a single CSP. Additionally, some CSPs offer the ability to deploy their services on-premises using pre-packed hardware and virtualization; therefore, an agency may have some CSP services running on-premises, in satellite or remote offices, in data centers, and/or in the cloud. 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; however, application or API connections to these environments must be secured by both agencies and the service providers.

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

In PaaS, vendors offer platforms, such as web servers and databases, to build solutions. Some PaaS features are often included as part of IaaS but can also be offered independently. 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.¹²

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. Agencies operating in a multi-cloud environment need to optimize their 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 are encouraged to use tools that provide a holistic view of their application and infrastructure across all CSPs to manage security policy in a centralized way. Agencies also have the choice to use tools that are offered by CSPs and by third-party vendors for security analysis across multiple CSPs. Agencies will want to determine which of these tools best improve their security posture based on their specific needs. Agencies should evaluate the benefits and shortcomings of security tools offered by CSPs and independent tools designed for multi-cloud environments. Where possible, agencies should use security tools that can work across multiple CSPs.

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

¹² National Institute of Standards and Technology, “NIST Special Publication 800-145: The NIST Definition of Cloud Computing,” (2011), <https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf>.

location such as authentication logs if using an integrated identity access provider across multiple CSPs. Agencies must be aware of and follow OMB Memorandum (M)-21-31 for log management.¹³

When planning to adopt cloud services agencies must determine how they will implement authentication and access management for each service. They must consider the implications associated with 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). Agencies should implement the strongest security features wherever possible such as implementing phishing-resistant multi-factor authentication (MFA)^{14,15}, and they should consider when to use convenience features like 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 deployment settings.

3.2 Introduction to FedRAMP

FedRAMP was established in 2011 by the OMB Memorandum, “*Security Authorization of Information Systems in Cloud Computing Environments*,” known as the FedRAMP Memo¹⁶. 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 government-wide program that promotes the adoption of secure cloud services across the Federal Government by providing a standardized approach to security and risk assessments for cloud technologies and federal agencies. As described in the FedRAMP Memo, FedRAMP is applicable to:

- Executive departments and agencies procuring commercial and non-commercial cloud services that are provided by information systems that support the operations and assets of the departments and agencies, including systems provided or managed by other departments or agencies, contractors, or other sources.
- All cloud deployment models (e.g., Public Clouds, Community Clouds, Private Clouds, and Hybrid Clouds) as defined by NIST.
- All cloud service models (e.g., Infrastructure as a Service, Platform as a Service, and Software as a Service) as defined by NIST.

The FedRAMP Memo further requires each Executive department or agency to:

¹³ “Improving the Federal Government’s Investigative and Remediation Capabilities Related to Cybersecurity Incidents,” Office of Management and Budget, (2021), <https://www.whitehouse.gov/wp-content/uploads/2021/08/M-21-31-Improving-the-Federal-Governments-Investigative-and-Remediation-Capabilities-Related-to-Cybersecurity-Incidents.pdf>.

¹⁴ Office of Management and Budget, “OMB M-22-09. Moving the U.S. Government Toward Zero Trust Cybersecurity Principles,” (2022), <https://www.whitehouse.gov/wp-content/uploads/2022/01/M-22-09.pdf>.

¹⁵ In this document, as in OMB M-22-09, “phishing-resistant” authentication refers to authentication processes designed to detect and prevent disclosure of authentication secrets and outputs to a website or application masquerading as a legitimate System.

¹⁶ Office of Management and Budget, “Security Authorization of Information Systems in Cloud Computing Environments,” (2011), https://www.fedramp.gov/assets/resources/documents/FedRAMP_Policy_Memo.pdf.

- Use FedRAMP when conducting risk assessments, security authorizations, and granting Authority to Operate (ATO) for all Executive department or agency use of cloud services.
- Use the FedRAMP Program Management Office (PMO) process and the Joint Authorization Board (JAB)-approved FedRAMP security authorization requirements as a baseline when initiating, reviewing, granting, and revoking security authorizations for cloud services.
- Ensure applicable contracts appropriately require CSPs to comply with FedRAMP security authorization requirements.
- Establish and implement an incident response and mitigation capability for security and privacy incidents for cloud services in accordance with DHS guidance.
- Ensure that acquisition requirements address maintaining FedRAMP security authorization requirements and that relevant contract provisions related to contractor reviews and inspections are included for CSPs.
- Require that CSPs route their traffic such that the service meets the requirements of the Trusted Internet Connections (TIC) program, consistent with DHS guidance.
- Provide, to the Federal Chief Information Officer (CIO) annually on April 30, (1) a certification in writing from the Executive department or agency CIO and Chief Financial Officer (CFO) and (2) a listing of all cloud services that an agency determines cannot meet the FedRAMP security authorization requirements with appropriate rationale and proposed resolutions.

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.
- Provide guidance for agencies and vendors to leverage for acquiring secure cloud solutions.

FedRAMP is continuing to look at ways to modernize and automate in service of our program mission. FedRAMP partnered with NIST and industry to develop the Open Security Control Assessment Language (OSCAL)¹⁷, a set of formats expressed in XML, JSON, and YAML. These formats provide machine-readable representations of control catalogs, control baselines, system security plans, and assessment plans and results. OSCAL is being applied to FedRAMP baselines and security package materials in order to streamline the development and review of authorization packages. To aid users in getting started with OSCAL, FedRAMP additionally released open source tooling, to include OSCAL Generator and Conversion tools¹⁸. To build upon the foundation established in Fiscal Year 2021, FedRAMP will continue to prioritize continuous improvement of business processes that will help all stakeholders. Benefits will impact key stakeholder groups in the following ways:

¹⁷ National Institute of Standards and Technology, “OSCAL: the Open Security Controls Assessment Language,” <https://pages.nist.gov/OSCAL/>.

¹⁸ General Service Agency, “FedRAMP Automation,” <https://github.com/GSA/fedramp-automation>.

- Agencies will have an improved view into risk management, resulting in better informed decision making while authorizing cloud service products, ultimately enabling their organizations to adopt new services faster.
- CSPs and Third Party Assessment Organizations (3PAOs) will have automated mechanisms to self-test, develop, submit, and remediate security packages, reducing the level of effort and timeline for authorizations. CSPs will additionally have automated channels to conduct continuous monitoring, resulting in faster resolutions for cybersecurity threats.
- FedRAMP will receive improved packages at the outset of an authorization lifecycle, resulting in fewer setbacks during the review process. Through automated formats, package reviews will be streamlined, less cumbersome on stakeholders, and result in faster decision making.

3.2.1 FedRAMP's Stakeholders: Roles and Responsibilities

Four stakeholder groups serve roles in FedRAMP—CSPs, 3PAOs, federal agencies, and the 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 and be committed to understanding FedRAMP, leveraging FedRAMP templates to maintain alignment to and compliance with the shared responsibility requirements established by FedRAMP. 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 PMO provides training, guidance, and advisory support to CSPs, helping them navigate the FedRAMP process and understand the requirements. CSPs providing CSOs for federal consumption should be committed to understanding FedRAMP and leverage FedRAMP templates to maintain alignment to and compliance with the shared responsibility requirements established by FedRAMP.

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 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. To establish a consistent approach to federal cloud adoption, agencies and CSOs are encouraged to receive FedRAMP training and to develop system-level security artifacts using FedRAMP templates. 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 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 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.¹⁹

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. Agencies using cloud environments remain responsible for monitoring portions of the environment that CSPs do not monitor, which is generally covered under separate authorizations (See Section 3.1 for how the layers of the cloud service models work with various roles and responsibilities).

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).²⁰ 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. FedRAMP provides additional continuous monitoring guidance documents, such as the *FedRAMP Guide for Multi-Agency Continuous Monitoring*²¹. FedRAMP strongly encourages agencies to leverage this guide in order to share the

¹⁹ The Federal Risk and Authorization Management Program, “Joint Authorization Board Charter,” (2018), https://www.fedramp.gov/assets/resources/documents/FedRAMP_Joint_Authorization_Board_Charter.pdf.

²⁰ The Federal Risk and Authorization Management Program, “FedRAMP Continuous Monitoring Strategy Guide,” (2018), https://www.fedramp.gov/assets/resources/documents/CSP_Continuous_Monitoring_Strategy_Guide.pdf.

²¹ The Federal Risk and Authorization Management Program, “Agency Guide for Multi-Agency Continuous Monitoring,” (2020), https://www.fedramp.gov/assets/resources/documents/Agency_Guide_for_Multi-Agency_Continuous_Monitoring.pdf.

responsibility of continuous monitoring, reduce the dependency of leveraging agencies on the initial authorizing agency, and collaborate with the CSP and other member agencies to ensure the cloud service continues to meet the member agencies' needs. Additionally, agencies should consider using the *FedRAMP Continuous Monitoring Performance Management Guide*²² to provide a consistent approach to managing the security posture of CSOs in the continuous monitoring phase. To facilitate efficiencies through automation and tooling, with the permission of the CSP, agencies may incorporate security artifacts from vendors into agency governance, risk, and compliance (GRC) capabilities to ensure cloud service security posture is visible to agency risk management framework (RMF) stakeholders and authorizing officials.

3.3.2 Incident Handling

The Federal Information Security Modernization Act of 2014 (FISMA),²³ at 44 U.S.C. § 3552(b)(2), defines 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 phase. Clear and timely incident communication to relevant stakeholders is a key aspect of continuous monitoring 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*²⁴ 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

NIST defines the Security Authorization Boundary as "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."²⁵ 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 CSO's internal services, components, and other devices along with connections to external services and systems. An authorization boundary diagram encompasses all technologies, external and internal services, and leveraged systems and accounts for all federal information, data, and metadata that a

²² "FedRAMP Continuous Monitoring Strategy Guide," The Federal Risk and Authorization Management Program, (2018), https://www.fedramp.gov/assets/resources/documents/CSP_Continuous_Monitoring_Strategy_Guide.pdf.

²³ Codified in relevant part at 44 U.S.C. § 3551, *et seq.*

²⁴ The Federal Risk and Authorization Management Program, "FedRAMP Incident Communications Procedure," (2021), https://www.fedramp.gov/assets/resources/documents/CSP_Incident_Communications_Procedures.pdf.

²⁵ National Institute of Standards and Technology, "Security Authorization Boundary," https://csrc.nist.gov/glossary/term/security_authorization_boundary.

CSP 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 document²⁶ 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 U.S./U.S. Territories or geographic locations where there is U.S. jurisdiction 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 U.S./U.S. Territories or geographic locations where there is U.S. jurisdiction concerns through SLAs or memorandums of understanding (MOUs).

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.
- **Centralize Common Cloud Services:** Identify CSPs that will be used across the agency and centralize the procurement and administration.
- **Invest in People:** Cloud migrations need specialized skills that agencies must cultivate.

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 be able to develop architectures that are scalable, repeatable, reliable, and align with zero trust philosophy. This process requires collaboration across agency teams to build digital services. DevSecOps can combine with centralized SaaS, supported by IT departments, to enable security testing of 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

²⁶ The Federal Risk and Authorization Management Program, “Requesting Public Comment on FedRAMP Authorization Boundary Guidance,” (2021), <https://www.fedramp.gov/blog/2021-07-14-Public-Comment-Boundary-Guidance/>.

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-premises 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. Agencies can initially transition smaller, internal projects and tools to the cloud to gain experience and confidence working in a new environment before attempting to migrate larger services. Shifting to cloud 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 architectures, and more secure applications, can be easier in the cloud. CSPs can address aspects of the five zero trust pillars—Identity, Devices, Networks, Applications, and Data—and enable the visibility needed to begin creating cross-pillar interactions²⁷. 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 ICAM 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; as data accumulates over time, it can pose additional challenges to cloud migration. When agencies decide to migrate their application

²⁷ Cybersecurity and Infrastructure Security Agency. "CISA Zero Trust Maturity Model," (2021), https://www.cisa.gov/sites/default/files/publications/CISA%20Zero%20Trust%20Maturity%20Model_Draft.pdf.

ecosystem to the cloud, they should weigh benefits, risks, and challenges to adopting cloud-based technologies.

4.2.1 Possible Cloud Migration Challenges

All large-scale software projects have their challenges but moving 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 2: 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 inexpensive or even free, depending on the CSP, the architecture, and the approach; moving data out can be more costly.
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 existing application/project ATOs, they may need to be updated or replaced by new ATOs.
Change Management	Moving to a cloud architecture will require changes in process, in addition to the technical changes. Acknowledging this and creating space to remake the processes will ease some of the discomfort of changing.

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 3: 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 via encryption 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.

Technical Challenges	How does it affect the mitigation?
Network Support ²⁸	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 decisions on how to better migrate the data to the cloud vendor's environment. Bandwidth may also be an issue for developers having to move data and applications around, as well as for end users on home networks.

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 800-145²⁹ 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. Tenants can opt for virtual machines (VMs) instead of reserving hardware. 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 some of the routine work of health monitoring and patch management to the CSP, though agencies would remain accountable for the security of their systems. 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 but is not all inclusive.

Table 4: 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. Scalability is key to distributed systems.
Availability	Cloud services can manage failures of the underlying infrastructure for the application so that running code can be moved with minimal interruption.
Cost	CSP services can increase efficiency while allowing agencies to direct financial resources towards mission-critical tasks.
Disaster Recovery and Business Continuity	Agencies with off-premises cloud data and infrastructure are better positioned to handle and recover from adverse events at agency offices (e.g., natural disasters).
Cybersecurity	CSPs often provide options for different aspects of security so individual customers do not have to build out their own support for it. However, it is crucial

²⁸ Transferring data over an agency's network is only one option. There may be other services that can be used to migrate data into the cloud, such as copying data to disks and transporting them to the CSP by ground or air.

²⁹ National Institute of Standards and Technology, "NIST Special Publication 800-145: The NIST Definition of Cloud Computing," (2011), <https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf>.

Benefits	How does it benefit a project?
	that agencies learn about the options and implement and configure the ones that are right for them.

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. Since not every application is designed to run in a cloud environment, agencies must consider their specific needs as they migrate. For example, an application may depend on the low latency provided by a local network and a CSP might not be able to provide that speed.

Table 5: 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.

There is much debate between the Rehost strategy and the Refactor or Revise strategies, and agencies should carefully consider which one is right for them. There are times when it is necessary to move an application to the cloud due to legacy system deprecation but attempting a Refactor at the same time is not feasible. In that case, the right strategy might be to pursue the Refactor after the Rehosting is complete. The Refactor should still be considered as there are many ways in which cloud native services from an IaaS or PaaS can reduce complexity, improve performance, and lower hosting costs.

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

Every cloud migration is as unique as the original application, thus it is challenging to give universal recommendations on how to perform the migration. However, following the phases below can increase the chances of success.

- **Plan:** Determine which strategy to use, which CSP and service type, and the road map for the application.
- **Design:** Create the architecture for the application focusing on the distributed nature of the system. Trial cloud-native features of the CSP for use.
- **Pilot:** Create a Minimum Viable Product (MVP) to demonstrate that the application will work in the cloud.
- **Migrate:** Make the cloud version production ready, including porting over any needed data.
- **Maintain:** Continue improving the cloud application, whether from a product feature perspective or from a performance perspective.

The following subsections outline common migration scenarios for agencies. As these scenarios are focused on the ways that application architecture changes when moving to a cloud environment, they do leave out the security functionality that is routine to the environment.

4.3.1 Scenario 1 – PDF Storage 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.

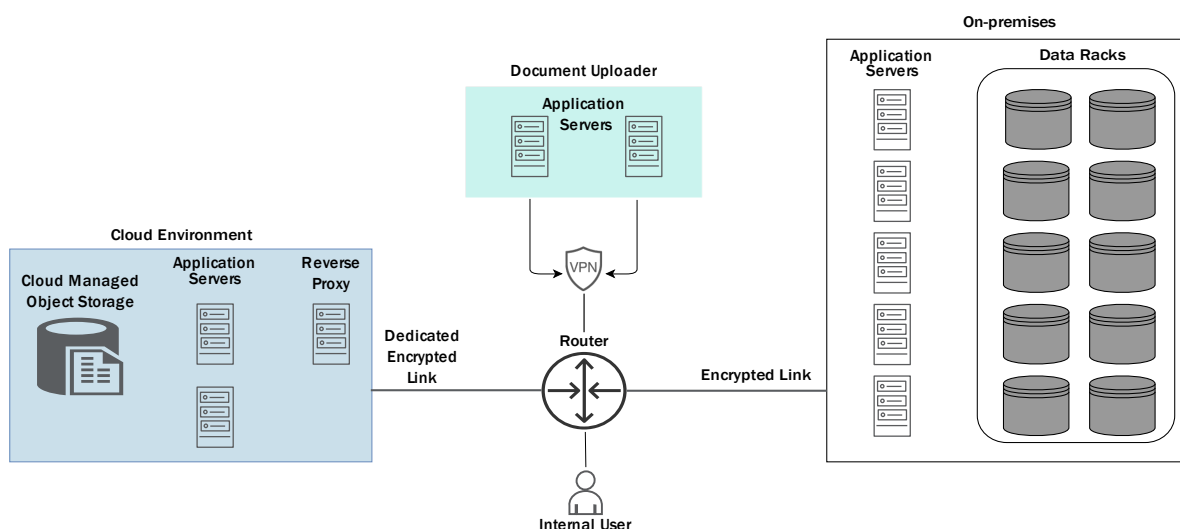


Figure 3: Scenario 1 – Notional Phase 1 Architecture

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 or other storage. This technique may be more efficient than transferring all the data over the network.

Figure 4 shows these adjustments.

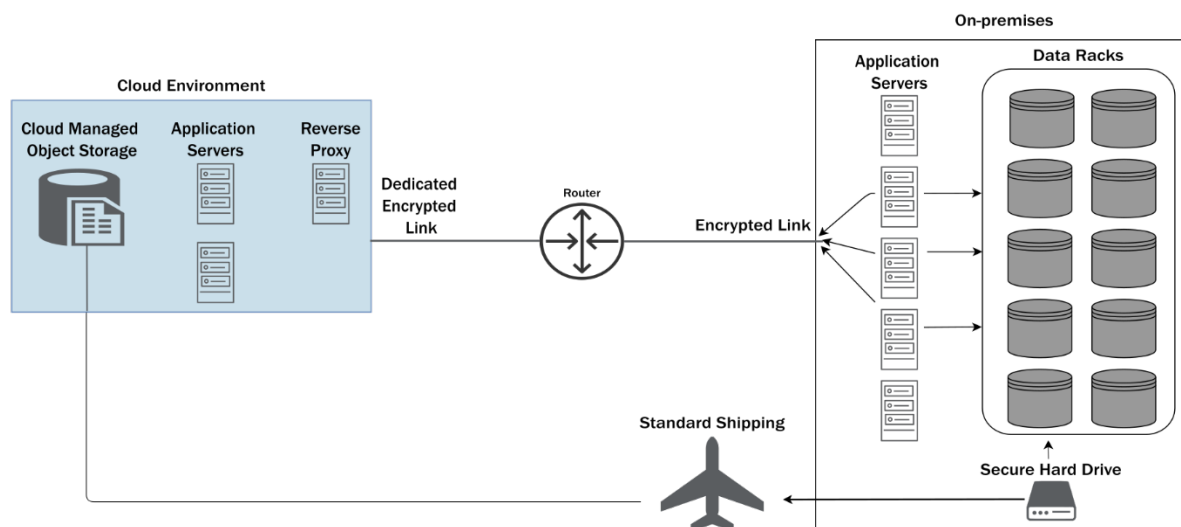


Figure 4: Scenario 1 – Phase 2 Notional Architecture with Out-of-Band Data Transfer

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, the agency 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.

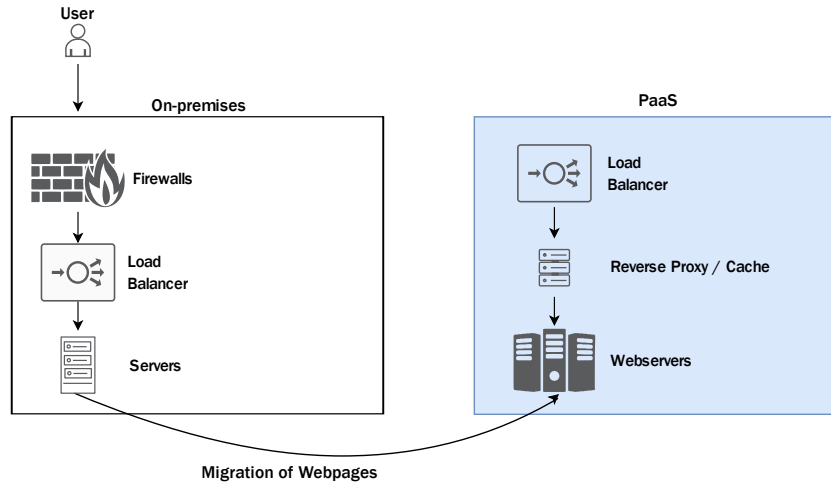


Figure 5: Scenario 2 – Notional Migration of a Website to a PaaS

After the migration and redesign, the agency recognized that most of the webpage content is public and does not change frequently and thus 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 of files to the CDN and configure it to serve user traffic. Agencies should assess the data to be cached in a CDN service. Many CDNs offer additional security features such as Distributed Denial-of-Service (DDoS) attack mitigations and web application firewalls (WAFs) that agencies can also take advantage of. Most data will be public, and therefore acceptable to be cached outside an authorized boundary. Some data will have CUI requirements, and therefore should be uncached, or the agency should use an authorized CDN provider. Figure 6 shows one example of migrating a website to PaaS.

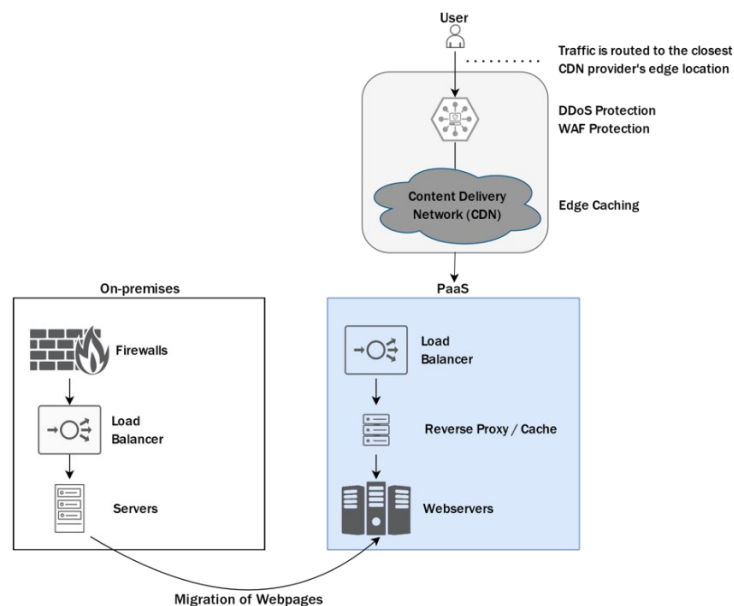


Figure 6: Scenario 2 – Notional Website with CDN

The desired 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.

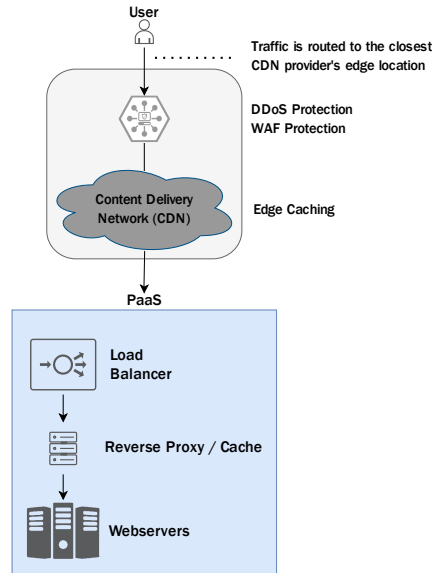


Figure 7: Scenario 2 – Notional Final Architecture of the New Website

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 performance 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 designed to generate the synthetic traffic and collect the resulting metrics. The team settles on using a SaaS system (Figure 8).

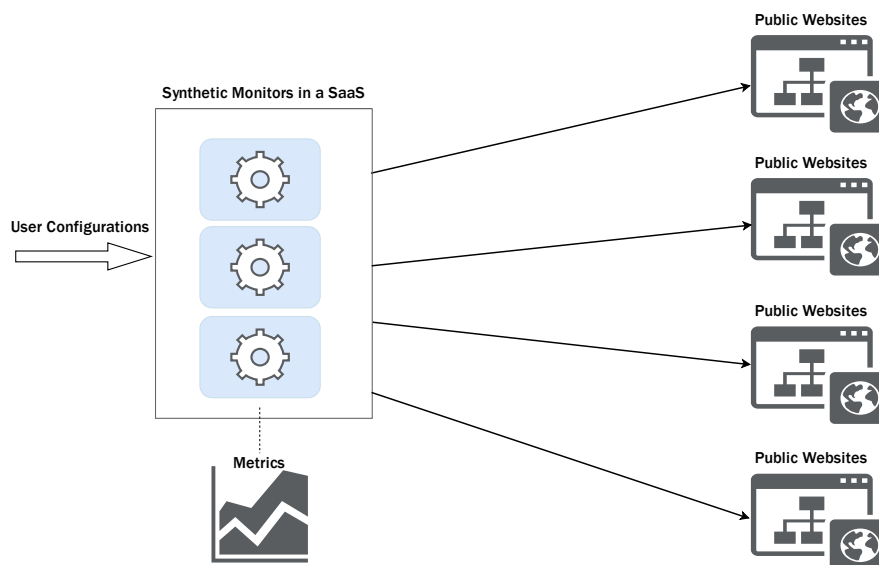


Figure 8: Scenario 3 – Notional Deployment of SaaS-based Website Monitoring

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 successfully 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), Infrastructure as Code (IaC), security testing, and the principle of least privilege to harness automation and produce reliable and predictable digital services that scale.

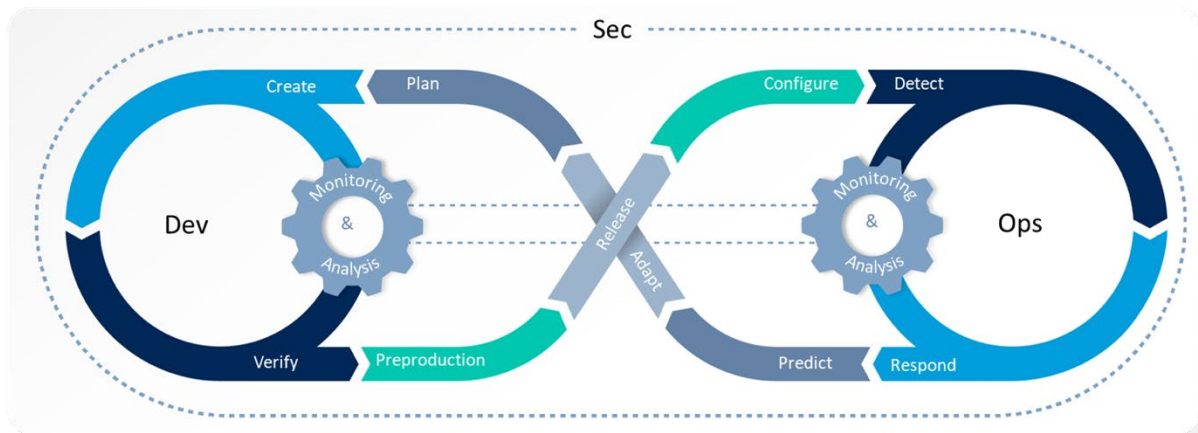


Figure 9: DevSecOps Loop

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. Source code management software can also enforce procedures for code review and code check in that further reduce human errors and add non-repudiation into the system³⁰.

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.

³⁰ National Institute of Standards and Technology, “SP 800- Secure Software Development Framework (SSDF) Version 1.1: Recommendations for Mitigating the Risk of Software Vulnerabilities,” (2021), <https://csrc.nist.gov/publications/detail/sp/800-218/draft>.

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 can produce more reliable software when it is time for deployment. Not only does this practice help 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 Infrastructure as Code (IaC) and it enables teams to review changes to the resources used in IaaS or PaaS before checking in that code. It also facilitates mass production of cloud infrastructure so patches can be applied quickly, and environments can auto-scale. Artisanal servers need to be individually patched; thus, the state of the infrastructure is prone to drifting away from the original configuration when manually updated.

IaC can offer a multitude of benefits:

- Removing the need for a User Interface (UI) at each device, which further reduces opportunities for human error;
- Automating compliance checks using IaaS or PaaS features, such as enforcing encryption at rest for storage containers;
- Automating deployment of ICAM policies as well as granular access controls;
- Facilitating security testing, patch deployments, and updates; and
- Increased zero trust maturity through enforcing encryption on networks and storage through code.

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 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 crucial way to integrate security within an earlier phase of the software development lifecycle. 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 and when they are easier to fix. Testing throughout the CI/CD process also leads to increased zero trust maturity.

Automated security testing during development is just one layer of defense against application vulnerabilities. The layers of manual expert analysis, third-party security testing, and public vulnerability disclosure programs along with bug bounty programs work together to ensure applications are exercised and increase the chance of vulnerabilities being identified before they can be exploited. See CISA

Binding Operational Directive 20-01³¹ 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 be 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.

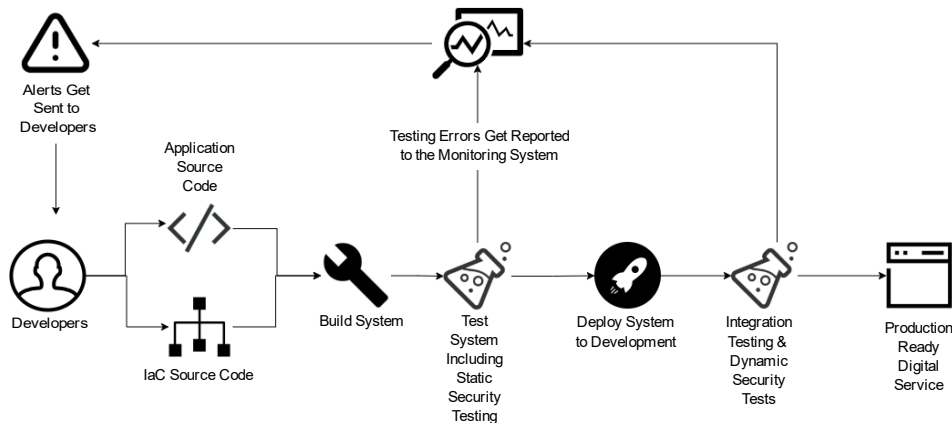


Figure 10: Reference Architecture for a Build System with Security Testing

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 member has sufficient privileges to do his or her job, but no more privileges than what that user needs. 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. This helps minimize the risk of misuse by a malicious actor (internal or external) by limiting how they can elevate privileges or restricting possible movement. Some 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 also critical.

³¹ Department of Homeland Security, “Binding Operational Directive 20-01: Develop and Publish a Vulnerability Disclosure Policy,” (2020), <https://cyber.dhs.gov/bod/20-01/>.

Related to least privilege is the rise of attribute-based access control (ABAC)³². ABAC takes role-based controls a step further by enforcing checks around the user's identity, the attributes of the resource being accessed, and the environment. Roles then become an attribute of the user's identity. Another equally important attribute to check for is, "May the data be accessed by this user?" Additionally, a common environment-based attribute check is information about the device the user is using—Is it an agency device that is up-to-date on its patches? Combining multiple attributes can give higher confidence that the user is who they say they are and that they are permitted to perform the requested action. ABAC is a core component of a mature zero trust architecture by involving more than one pillar in access decisions.

Traditionally, separation of duties has been used to deter insider threats and catch innocent mistakes by requiring more than one individual to perform important tasks. An example is the team that does development and coding is separate from the team that does production deployment. This approach is in tension with DevSecOps since these responsibilities are now shared within a team.

A replacement process is a two-person integrity check approach through code reviews. This means every code and configuration change submission must be reviewed and approved by another authorized team member before the change is committed and merged into the main repository. This is useful in both the application code as well as IaC to catch issues before deployment. Many code repositories can be configured to enforce code reviews. Additionally, the repository administrator accounts that allow this setting to be disabled should not be used for regular, daily activity.

4.5 Centralizing Common Cloud Services

As developers migrate, create, and deploy applications in the cloud, their agency can help by managing and maintaining shared services. By providing shared services, agencies allow developers to spend more time focused on the mission and less time on overhead or maintenance tasks. These services are broken into four areas here:

- Agency PaaS,
- Development tools and services,
- Public-facing services, and
- Security services.

Sharing some services at an agency level can help teams begin using cloud native techniques faster by removing administration overhead, leading to the freedom to think about other overhead that can be removed. A team can move away from running full VMs for web servers to running servers out of containers and then moving from containers to using "Functions-as-a-Service"³³.

There is another evolution as agencies move from traditional on-premises servers to IaaS in that specialized roles are used differently. Teams will no longer need someone dedicated to administrating the server that the database is on but will still need the specialized knowledge to understand how to make

³² NIST defines ABAC as "An access control method where subject requests to perform operations on objects are granted or denied based on assigned attributes of the subject, assigned attributes of the object, environment conditions, and a set of policies that are specified in terms of those attributes and conditions."

³³ Function as a service (FaaS) is a category of cloud computing services that provides a platform allowing customers to develop, run, and manage application functionalities without the complexity of building and maintaining the infrastructure typically associated with developing and launching an app.

databases performant. This model allows database administrators can focus on the strategic operations as well as fulfill a consulting role within the agency.

4.5.1 Why Centralize?

There are two main reasons to centralize IaaS, PaaS, and SaaS across an agency: to save resources and to build a shared experience. Procurement of new tools and software is resource intensive. Agencies can reduce overall cost by creating a centralized team that is responsible for research, procurement, and training on tools that all teams will use. In addition, centralizing services (like those in the subsequent sections) can also conserve resources by streamlining maintenance and compliance efforts.

Centralization also allows agencies to build a shared experience by making common tools available to different teams within an agency, enabling collaboration. Centralized documentation enables knowledge to spread outside of a team. Using the same ticketing, pager, and monitoring helps teams work together when there is an outage of cloud service. It also facilitates onboarding as employees transfer between teams in an agency.

Teams that have experienced cloud-based projects can also share their best practices and challenge areas so that other teams can learn from their experience. Members of experienced teams can mentor newer teams to share the knowledge and skills gained, increasing the overall investment in people. Reducing resources and breaking down silos in organizations are two strong reasons to centralize cloud tools.

4.5.2 Agency Platform-as-a-Service

Agencies can centralize access to current IaaS tools by procuring cloud infrastructure in bulk and provisioning access to different teams as needed. This will ensure the appropriate level of access is granted and will also allow newer teams to begin using the infrastructure quickly. An agency's cloud team can act more like a PaaS by offering configurations to standardize operating systems, software libraries, and logging. Together, these principles will accelerate the development of digital services in the cloud and save resources.

A centralized IaaS can establish normative behaviors and enforce compliance while beginning to reduce the burden of security paperwork like ATOs on development teams. Major IaaS platforms enable compliance checks, such as notifying teams when a storage container is public or not encrypted, so teams can fix the issue quickly. When all the teams share the same platform, they can inherit NIST SP 800-53³⁴ controls from the organizational account and use common language in their software support program (SSP) agreements for faster paperwork.

Agencies can centralize “gold image” VMs and establish artifact repositories so that teams can share containers used in the IaC. The VMs and containers can also be set up with the logging standards outlined in OMB M-21-31. The tension between usability and security appears here – while agencies can add security monitoring to the base images, it is important to also keep the images performant by not overloading the systems with too much extra processing. Additional security gains can be made by enforcing regular patching throughout the environment.

³⁴ National Institute of Standards and Technology, “NIST Risk Management Framework,” <https://csrc.nist.gov/Projects/risk-management/sp800-53-controls/release-search#!/800-53>.

Artifact repositories, along with IaC, move development teams closer to the idea of “immutable workloads;” once cloud infrastructure and code is deployed, it is not manually upgraded or changed. Any changes would be done through the CI/CD pipelines and the systems would be re-deployed.

Encryption services can ensure the application uses secure communication channels (such as TLS) by providing certificates for the application on the web server. Also, where necessary, these services can encrypt data at rest either directly or through managed services.

Key and password management enables applications to rotate keys and passwords on a timed basis without interruption to the application. This service must also have the capability to revoke keys if compromised.

4.5.3 Development Tools and Services

Development tools and services are key to quickly and efficiently building and maintaining applications. This section includes common tools and services used in application development but is not exhaustive.

The software development lifecycle, along with DevSecOps, uses collaboration tools, requirement tracking, and documentation extensively to share current state assessments within teams and across teams. Agency-wide collaboration and documentation practices create cultures of sharing and collaboration.

The source control product is the foundation of a CI/CD pipeline, as it drives what tools can be used for building, testing, and deploying code. Performing code quality control in the form of “linters,” which can examine code for issues that would prevent execution or create difficult to read code, and checking for coding “anti-patterns,” which can prevent poor coding conventions that create insecure or non-optimized code, is also an important part of CI/CD that can be standardized across the agency.

Security testing that can be integrated into the CI/CD pipeline is also important to centralize and standardize across an agency because the uniform application of security makes for better overall processes. Static and dynamic security testing can provide an early layer of defense from accidentally deploying bugs to production.

4.5.4 Public-facing Services

Some aspects of digital services that agencies provide to the public would benefit from centralization.

Routine aspects of deploying a new website include obtaining a domain, configuring domain name system (DNS) entries for the site, and setting up certificates for hypertext transfer protocol secure (HTTPS). Centralizing these processes also helps the agency maintain an accurate inventory of their web presence.

Applications and APIs that are accessible from the internet need protection from malicious traffic. Protection can come in the form of WAFs, API gateways, and content delivery networks (CDNs) that double as DDoS protection. WAFs can control access to the network in general as well as inspect the requests to the web server to look for common website attacks. An API gateway controls access to APIs for specific users, and multiple APIs can be protected by the same gateway.

CDNs not only provide a way to store cached data closer to users for faster delivery, but they often can also absorb extra traffic in denial of service (DoS) attacks or limit network traffic via firewalls. All internet properties of an agency will need this protection, and bulk purchasing can offer cost savings.

4.5.5 Security Services

Agencies should deploy centrally integrated security services to the greatest extent possible across the enterprise. Fewer separate instances of the same service reduce an agency's attack surface. Security services provide application protections such as logging, authentication, authorization, encryption, and key management.

Centralized logging is key to better incident response. It makes locally stored logs redundant and reduces the impact of their deletion. Centralized logging also reduces the amount of time needed for an incident responder to investigate. OMB M-21-31 also makes provisions for components to share logs with their parent Department, and centralization streamlines this process. Centralized logging also facilitates threat hunting across the CSPs and on-prem solutions.

For agency services, ICAM 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.

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 ICAM services to the cloud as the primary identity provider.

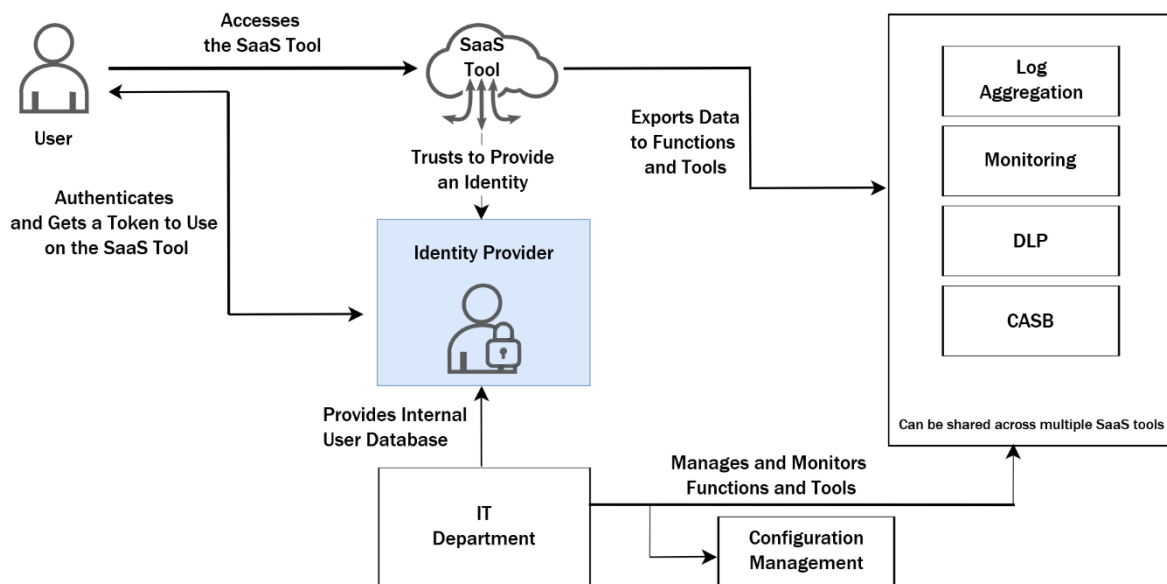


Figure 11: Reference Architecture on Centralized Security Services

Agencies should strive for minimal friction when onboarding a new SaaS product to the centralized ICAM system. They should explore opportunities to prototype or pilot connecting such services prior to full adoption and integration. This can allow agencies to address initial performance and security related concerns and provide additional insight for improved integration.

Authorization (i.e., the process of determining if a user has permission to perform an action) is often included in ICAM services. However, sometimes authorization needs to be enforced by applications. Development teams need to be able to continually request information about authorization. As agencies continue their zero trust journey, authorization will move from role-based to attribute-based and will need to encompass information from multiple pillars. Centralization also benefits authorization.

4.6 The Human Element

Building scalable, repeatable architectures via CSPs requires changes to processes and procedures, not only for the staff working on deploying tools and applications, but also for the stakeholders and users of these tools. Agencies will need to invest in people to deliver cloud-based projects. They will also need to redesign processes, educate all staff, and facilitate reliable access.

4.6.1 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). 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 way 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)³⁵. 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 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 Hub³⁶ 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)³⁷ to help procurement professionals learn more.

³⁵ President's Management Agenda, "How to Hire the Best Talent Across Government," (2020), <https://www.performance.gov/cx/blog/CX-hiring-pilot/>.

³⁶ United States Digital Service, "TechFAR Hub," <https://techfarhub.cio.gov/>.

³⁷ United States Digital Service, "Digital IT Acquisition Professional Training (DITAP)," <https://techfarhub.cio.gov/initiatives/ditap/>.

4.6.2 Support Staff

Through onboarding and other documented procedures, agencies should strive to support federal staff through additional training and timely access. All personnel will require some additional training, both on the use of new CSP tools and how the use of cloud tools changes the security paradigm. Security training may include anti-phishing training and proper data handling.

This new security paradigm will require adjustment from both personnel directly involved in the creation of digital services (e.g., those working in DevSecOps), as well as those supporting digital services from a non-technical perspective. Improved communication between development teams and stakeholders can help break down silos that may have grown over time. Working across divisions and mixing development, security, and operations will encourage collaboration and dismantle information silos.

Providing timely access also reduces the likelihood that employees will develop "shadow IT" services that circumvent oversight by IT or security teams and weaken the Federal Government's overall cybersecurity posture.

5. Cloud Security Posture Management

This section introduces cloud security posture management (CSPM) and the related security capabilities and outcomes. This section also highlights some key considerations when migrating to the 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.

The following 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 toward desired security outcomes.
- **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 means a continuous process of monitoring a cloud environment by identifying, alerting on, and mitigating cloud vulnerabilities; reducing risk; 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.

In this document, CSPM capabilities seek to support the following activity outcomes:

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

These capabilities include:

- Security and Risk Assessments,
- Continuous Monitoring and Alerting,
- Identity, Credential, and Access Management (ICAM),
- DevSecOps Integration, and
- Artificial Intelligence (AI)- and Machine Learning (ML)-Based Security Capabilities.

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 offload physical access to these deployed resources and change how they manage governance and compliance requirements for their applications and data. As 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 offered by cloud services are necessary for defending the Federal Government from cyber threats. Agencies should manage their risk 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, which enable agencies to keep up with emerging threats, protect against misconfigurations, and reduce the risk of a security incident or data breach. While some Agencies may be better poised to take advantage of these capabilities, preparatory activities, such as developing a warm site, see Appendix A – Scenarios, may offer all agencies immediate security benefits, operational resilience, and a foundation to adopt further capabilities.

5.1.2 How can CSPM facilitate Zero Trust?

As described by Executive Order 14028, agencies migrating to cloud deployments should adopt zero trust principles and transition their environments to zero trust architectures³⁸, as practicable, 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, application security, and data protections—integrated across environments on-premises and in the cloud. Additionally, agencies should apply automation and orchestration, governance, and visibility across these areas. As noted in Section 3.1 of NIST SP 800-207, there are several ways that agencies can design a zero trust architecture, including via enhanced identity governance, logical micro-segmentation, and network-based segmentation approaches. However, a full zero trust solution will include elements of all three of those approaches. Agencies can also use CISA's Zero Trust Maturity Model³⁹ for developing a strategy to adopt Zero Trust.

First and foremost, agencies should work towards an identity management solution that provides enterprise-wide 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 associated credentials as well as 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, including for service, network, and workload identities, to provide monitoring and analysis and ensure access controls are automatically configured for deployed services at scale and across environments.

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

In a zero trust architecture both on-premises and in cloud environments, agencies should segment their networks to reduce lateral movement, limit permissions, and control attack vectors.⁴⁰ Agencies should

³⁸ National Institute of Standards and Technology, “NIST Special Publication 800-207: Zero Trust Architecture,” (2020), <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf>.

³⁹ Cybersecurity and Infrastructure Security Agency. “CISA Zero Trust Maturity Model,” (2021) https://www.cisa.gov/sites/default/files/publications/CISA%20Zero%20Trust%20Maturity%20Model_Draft.pdf.

⁴⁰ This is not intended to supersede other zero trust architecture approaches driven by enhanced identity governance or logical micro-segmentation, but rather function in concert as part of a full zero trust solution.

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 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 reassess how they secure their data in the cloud. Agencies should always protect data at rest in the cloud and in transit to, from, and within cloud deployments. 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,
- Standards and Policies,
- 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

In the process of developing and implementing cloud governance for guiding operations and deployment, agencies must comply with both 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, 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 requirements, offering a minimal level of compliance. Many of the solutions also include tools for the continual assessment of cloud deployments and environments against these requirements. As service providers implement changes and update terms of service, agencies should consider how those changes may natively support compliance with applicable requirements or may result in non-compliance and require additional remediation.

5.2.2 Standards and Policies

Beyond governance and compliance, agencies should consider industry standards and best practices to help 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 ICAM, data protection, and others.

5.2.3 Identity, Credential, and Access Management

One particularly important component of policy enforcement is the handling of identities, credentials, and access controls. CSPM tools can help integrate ICAM 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 ICAM controls, CSPM capabilities help ensure that services automatically inherit the appropriate configurations. 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 ICAM with granular account access controls, directory services, application and resource authorization, and policy compliance. With these built-in capabilities, the CSP then acts as an enabler for agencies in adopting a zero-trust approach, taking advantage of commercially adopted standards built on top of scalable infrastructure. Common cloud solutions also help to promote interoperability, efficiency and reuse, and federated access.

Agencies should ensure that their own on-premises ICAM controls are up-to-date and consistent with their CSPs' controls. Best practices such as enabling phishing-resistant 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.⁴¹ ICAM also includes the use of analytics for monitoring, along with auditing and reporting, to support compliance. Agencies can

⁴¹ Agencies should consult OMB M-22-09, "Moving the U.S. Government Toward Zero Trust Cybersecurity Principles," Section III A, for additional details on the use of MFA.

reference other architecture and governance guidance, such as from the Federal Identity, Credential, and Access Management (FICAM) program⁴² managed by GSA.

5.2.4 Data Protection

Agencies must work with their CSPs to determine responsibilities for data management and protection. Data protections are necessary 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, (2) how much protection to apply, and (3) 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 policies, laws, regulations, and standards to establish appropriate governance for data protection. Agencies should protect intellectual property with digital rights management, as appropriate, as well as properly configure cryptographic services such as key management and PKI or symmetric encryption. Re-assessments should also 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 5.3.10). 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. Emerging technologies for secure computation and operations may also be relevant for supporting data protection.

⁴² General Service Agency, "FICAM Playbooks," <https://playbooks.idmanagement.gov/>.

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, containers⁴³, 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,
- Provide context on service configurations, and
- Support prioritization of proactive responses through alignment of resources and applications.

Such security management and risk management dashboards and tools allow CSPs to take measurements and compile reports about the effectiveness of decisions.

These dashboards, tools, and reports may help 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 ensures that networking permissions, segmentation, firewall, proxy, certificates, etc., are set correctly to support secure use. Host, firewall, and other policies should support the ability to isolate systems based on attributes such as location, application, environment (e.g., development or production), or resource type. on attributes such as location, application, environment (e.g., development or production), or resource type. The ability to deploy different security policies for each service type enhances the overall security posture. Additionally, agencies should take precautions regarding network access and network security settings, for example encrypting connections, using phishing-resistant multi-factor authentication, etc. This aligns with the zero trust tenet in which all communication is secured.

Resource Protections

Resource protection, including CSP service configuration protection, is another key component of infrastructure protection. Resources in this context include any resource or service that a CSP offers that is used by a tenant. 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:

⁴³ While short-lived containers may not require scanning during execution, container images should be scanned for vulnerabilities in the pre-deployment phase.

- Using application layer firewalls,
- Implementing mitigations for distributed denial of service attacks,
- 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 includes exercising caution in how accounts within applications are managed, what permissions or access the accounts have, who has access to the accounts, and how they are protected. By conducting this identification and mitigation, agencies can protect the applications and enable quicker responses to their potential misuse. Strong application security is another 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 systems are kept 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. It should also include an external-facing aspect, namely a vulnerability disclosure policy (VDP), as specified by CISA's Binding Operational Directive 20-01⁴⁴.

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 security events and may trigger notifications to the users or automate 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 wide 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 use different data to indicate system health.

⁴⁴ Department of Homeland Security, "Binding Operational Directive 20-01: Develop and Publish a Vulnerability Disclosure Policy," (2020), <https://cyber.dhs.gov/bod/20-01/>.

5.2.7 Incident Response and Recovery

Agencies should establish and maintain plans to respond to and recover from cybersecurity incidents.⁴⁵ 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.⁴⁶ 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 to facilitate response and recovery, including robust backup policies and procedures⁴⁷, 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

⁴⁵ Cybersecurity and Infrastructure Security Agency, "Cybersecurity Incident and Vulnerability Response Playbooks," (2021), https://www.cisa.gov/sites/default/files/publications/Federal_Government_Cybersecurity_Incident_and_Vulnerability_Response_Playbooks_508C.pdf.

⁴⁶ National Institute of Standards and Technology, "NIST Cloud Computing Forensic Science Challenges," (2020), <https://csrc.nist.gov/publications/detail/nistir/8006/final>.

⁴⁷ National Institute of Standards and Technology. "NIST SP 1800-11 Draft Data Integrity Recovering from Ransomware and Other Destructive Events," (2017), <https://www.nccoe.nist.gov/publication/1800-11/VoIB/>.

ways in which security tools can be deployed independently or as part of an integrated deployment to support delivery of CSPM capabilities.

This includes:

- CSPM Capabilities,
- Independent and Integrated Capabilities,
- CSP Account Management Hierarchies,
- Identity, Credential, and Access Management,
- Evolution of the Perimeter,
- Visibility and Sensor Placement,
- Monitoring,
- Application Programming Interfaces,
- 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 solutions into their services, other CSPs 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, providing continuous monitoring and visibility into identities and their permission sets with an automated risk-based context. CSPs offer capabilities like traditional on-premises port, service, and configuration scans, which may be used to moderate accounts, inspect traffic, identify vulnerabilities in services, and analyze code repositories. 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 (e.g., supplementing continuous monitoring and alerting with the generation and analysis of network metadata, sourced pervasively from inside cloud environments). 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 (SIEM) systems can be used to collect, monitor, and alert based upon the logs provided by a CSP.
3. **ICAM 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. Agencies should not assume such services are secure 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
 - Executing real-time health and performance monitoring of assets;
 - Remediating misconfigurations from both users and automated deployments;

- Redeploying existing infrastructure in response to incidents;
 - Monitoring and redirecting traffic through CDNs to extend data visibility and access controls beyond the traditional network perimeter;
 - Segmenting networks and provisioning segmentation for workloads, container, and cloud objects; and
 - Using infrastructure as code (IaC) 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. These capabilities can improve analytics and automation with insight-based prioritization, though iterative review should be conducted to reduce the risk of both natural bias and adversarial machine learning.

While CSPs may provide many of the above services, agencies may also look to third-party solutions 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 solutions via capabilities like cloud access security brokers (CASBs), secure access service edge (SASE), and SECaaS offerings. These offerings can provide agencies the ability to outsource a portion of their security and monitoring responsibilities to a CSP or third party using automated, managed security solutions.⁴⁸ While combinations of traditional capabilities may achieve the outcomes described in Section 5.2, CSPM capabilities, either natively offered by a CSP or via a third party, facilitate these outcomes.⁴⁹

Table 6: CSPM Outcomes

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

⁴⁸ While CASBs, SASEs, and SECaaS are typically third-party, CSPs may also offer them natively.

⁴⁹ 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.

CSPM Outcomes	Security and Risk Assessments	Continuous Monitoring and Alerting	Identity, Credential, and Access Management	DevSecOps Integration	AI and ML-based Security Capabilities
Incident Response and Recovery	X	X	X	X	X

Agencies should consider how integrated they want to be with each CSP they use, as some capabilities may promote vendor lock-in. 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 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., stand-alone, or non-integrated capabilities) and/or integrated capabilities across cloud deployments to better identify existing vulnerabilities and on-going compromises, and to prevent future breaches. 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 develop their own solutions.

Alternatively, deployments may also be wholly based on integrated capabilities to handle unified coordination across services, such as 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 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 theoretical action flows:

- Validated (captured by the capability),
- Unvalidated (not yet captured by the capability), and
- Unmonitored (not captured by a capability).

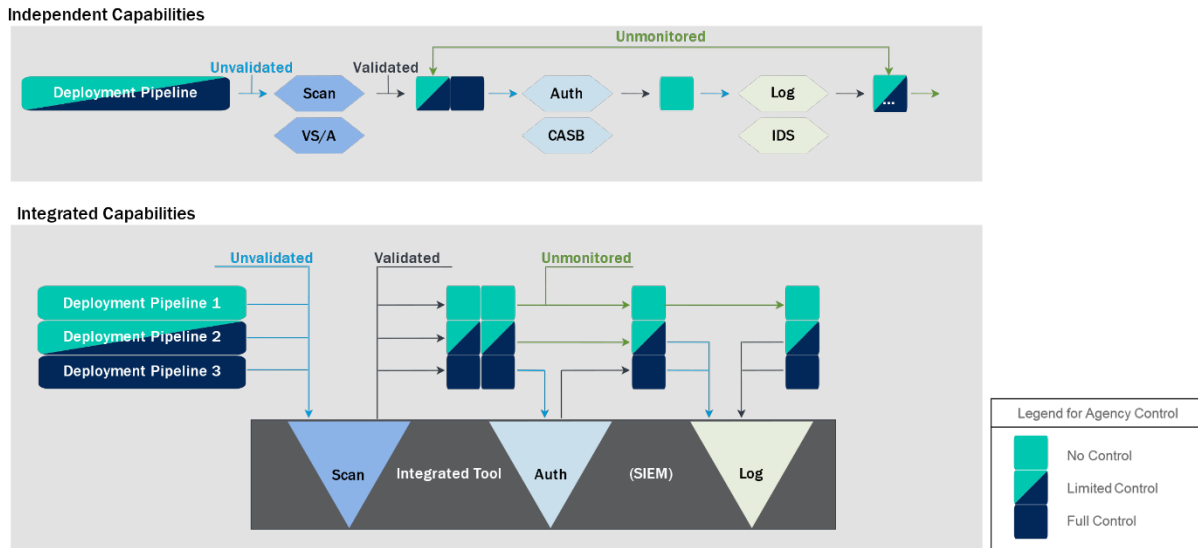


Figure 12: Service Deployments and Integrated Solutions

5.3.3 CSP Account Management Hierarchies

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. Naturally, **ICAM** capabilities are best suited to manage CSP Accounts, including the ability to audit activity, create and apply security rules, and set expirations on account validity. **Security and Risk Assessments** and **Continuous Monitoring and Alerting** capabilities may also be part of the auditing process both for user activity in real-time and for ensuring the safety of security rules.

Agencies may face other considerations regarding account organization and structure. Agencies should be aware that CSPs offering both commercial and government cloud services may not have the ability to move data internally between the two realms. This means that if an agency has a commercial account and a government account with a CSP and the agency wants all log information for accounts in a single location, then the agency would be required to pay to have the log data leave one of the accounts (either commercial or government) to 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, Credential, and Access Management

Identity and Credential 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. Agencies should refer to documents and resources such as NIST SP 800-63, OMB M-22-09⁵⁰, and the FICAM Playbooks to aid in both governance and compliance and to provide guidance on policies and procedures for identity and access management systems. 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. Additional details on federated identity can be found in the Federated Identity scenario in Appendix A. Some authentication services can integrate MFA and/or single sign-on. However, while many authentication providers may offer MFA, the MFA may not meet requirements for government systems, like PIV-enabled- or phishing-resistant-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.

CSP-provided **ICAM** capabilities can improve existing deployments by transitioning user management into a unified environment between on-premises and multiple cloud infrastructures. The **Security Risk and Assessment** and ICAM logs may be integrated using a SIEM to strengthen **Monitoring and Alerting** capabilities, and alerts may be created to trigger on specified users performing select activities. By enabling **DevSecOps Integration** and **AI/ML** capabilities to act on these alerts, agencies can:

- Automatically correct anomalous user behavior both in the CI/CD pipeline and via rollbacks in IaC, and
- Integrate access control into their service pipeline and apply behavior analysis to limit user's actions.

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.

⁵⁰ OMB M-22-09 Section III A includes the action, “Agencies must employ centralized identity management systems for agency users that can be integrated into applications and common platforms,” and includes additional details on enterprise-wide identity systems.

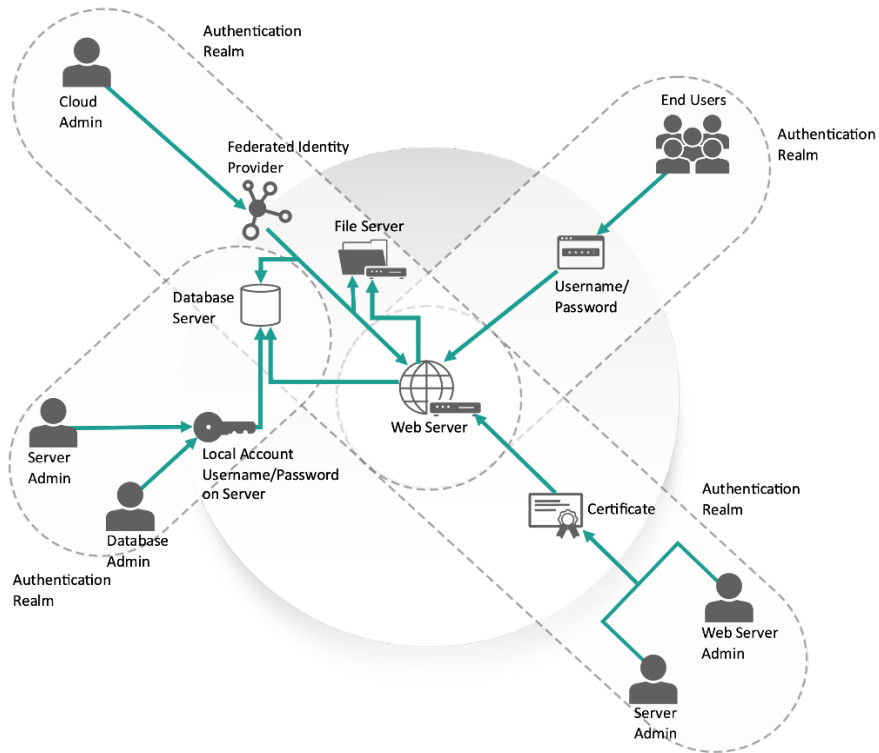


Figure 13: Authentication Realms

Agencies seeking to limit the number of authentications realms, as notionally shown in Figure 13, can use PaaS infrastructure, which eliminates the underlying servers that hosts the database and web server. The web server and database administrators would instead authenticate through the federated identity provider thereby eliminating the ways authentication can occur to system resources as seen in Figure 14.

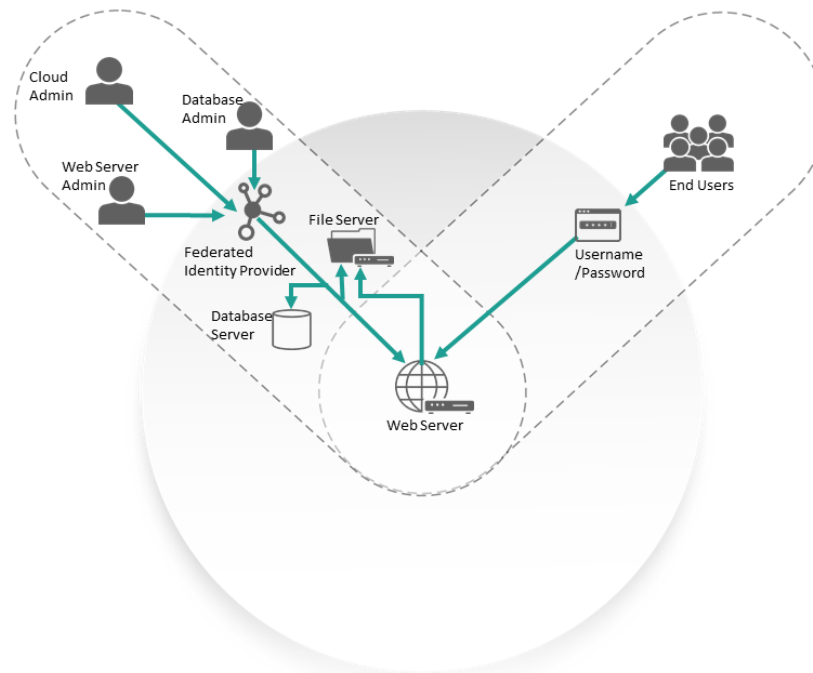


Figure 14: PaaS Authentication Example

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. Agencies may use auditing procedures in **Security and Risk Assessments** to detect over-privileged accounts and other account misconfigurations.

Because of the complex and dynamic nature of the cloud, agencies should consider implementing cloud infrastructure entitlement management as a way to enforce least privilege and to monitor authorizations for identities.

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. In IaaS environments, agencies will probably be able to emulate traditional network defenses and add action-based defenses that provide adversary detection through misinformation and redirection; such protections are likely unavailable in a SaaS environment.

Administrative access to a CSP's console (e.g., via either the web or command line) is open to the internet and agencies may 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.

Agencies should consider implementing a data-centric approach as part of their perimeter evolution, whereby security controls prioritize protection of data and assets before that of applications and services. Additionally, data, applications, and services can be segmented within the network to enforce more granular security policies for access to these resources.

By using security scans within **Security and Risk Assessment** capabilities and auditing within **ICAM** capabilities, agencies can analyze both their new cloud networks and test legacy networks for emerging vulnerabilities as they move to cloud. Because the perimeter is no longer limited to systems owned, managed, or used by agencies, the introduction of CDNs within **DevSecOps Integration** both increases the attack surface for data in transit and extends visibility of operations to the geographic region. To analyze data in transit through an extended perimeter, **AI/ML** capabilities may be used to better identify both classification and sensitivity of information to detect, predict, and track data exfiltration from both agency and CSP infrastructure. Additionally, **AI/ML** may protect the availability of services through adaptive load balancing and automated firewall management. Agencies may integrate all previously mentioned capabilities into their **Continuous Monitoring and Alerting** capabilities, improving awareness of existing and novel threats in both their existing and CSP-based infrastructure.

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 network 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 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 may not inform the tenant of actions taken on malicious or suspected malicious traffic as a result of those protections. Limitations on sensor placement and situational awareness vary by service model, offerings, and service configurations, which can impact the ability of the agency to act against threat actors and meet their visibility needs.

Agencies should consider sensor placement for both inbound and outbound traffic and monitor and control traffic between services within a network. All inbound and outbound traffic should be monitored as it enters and exits the tenant. Furthermore, agencies should consider monitoring traffic between services, especially traffic that moves through peered networks. Unauthorized or suspicious traffic should raise alerts and may be redirected to contained observation environments. With infrastructure at scale, the virtualization of interface monitors, network traffic sensors, and system logging **Continuous Monitoring and Alerting** capabilities may be leveraged to improve agencies' understanding of their cloud-based infrastructure. Maintaining control over the CDN as part of the **DevSecOps Integration** capability can also enable agencies to extend visibility with sensors placed well outside the agency perimeter, gaining further perspective over user and adversarial traffic.

5.3.7 Monitoring

Agencies may achieve a robust monitoring capability through the combination of CSPM capabilities, including **Continuous Monitoring and Alerting** and **DevSecOps**. Such capabilities can verify compliance, scan vulnerabilities, verify system availability under normal operations and simulated conditions, identify misconfigurations, and facilitate remediation. More specifically, these monitoring capabilities can identify and enumerate service uptime, quality of service, mitigate malicious behavior by ensuring content integrity, and help analysts and engineers investigate compromises and maintain uptime during incidents. See Section 5.3.9 for additional information on logging.

Agencies can use **Continuous Monitoring and Alerting** capabilities to augment their monitoring solution(s), like a SIEM, and enable scalability for both emerging and legacy log aggregation systems. Within monitoring applications, agencies can combine automated performance improvements and analytics from **AI/ML** with their chosen authentication deployment provider, from **ICAM**, to improve quality of service, content integrity, and service uptime requirements. As with the evolution of the perimeter, visibility, and sensor placement, CDNs within **DevSecOps Integration** enable agencies to gain a deeper understanding of underlying network constraints—such as regional availability, quality of service, and demand—that may hinder the ability to provide services to users. CSPs and third parties provide dashboards and additional tools allowing agencies to detect misconfigurations across their environments by analyzing IaC, service configuration, and ICAM permissions.

Monitoring should be used by agencies to track the footprint of the services they are using. This can serve two major purposes:

- Maintain and manage inventory of CSPs, CSP region operations, services, applications, accounts, and other assets; and
- Detect unauthorized use of services (e.g., shadow IT) that may occur by staff who operate their own accounts with CSPs and/or operate in unapproved regions.

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 aggregate 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. To accomplish this, monitoring services should ensure patches are properly applied and report the state of deployments across the cloud. 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. APIs 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 employ automation and efficient controls, apply best practices that minimize environmental drift, and enable the use of their services by third parties. As networks grow, agencies will need to responsibly manage the complexity associated with scaling services from both user and backend infrastructure. Agencies are encouraged to consult with NIST's special publications on security strategies for using microservices, i.e., NIST SP 800-204 and parts A, B, and C,^{51,52,53,54} as these resources informed the recommendations in this subsection.

By integrating CSP-based **Continuous Monitoring and Alerting**, agencies can collect real-time information about their usage of CSP-based APIs and other service usage data. In addition, **ICAM** capabilities can be integrated into API services to limit and control access, ensuring agency least privilege policies are applied. CSPM-based **Security and Risk Assessment** capabilities may be used to ensure agencies employ APIs for the management of their cloud-based infrastructure while maintaining compliance with privilege monitoring and vulnerability assessments. API applications and the CI/CD pipeline within **DevSecOps Integration** are mutually beneficial: development within the CI/CD pipeline ensures proper validation is conducted in the usage of APIs and using APIs within the CI/CD pipeline streamlines service patches.

Agencies may increase their attack surface by adopting APIs because they include externally developed code for which agencies have neither control nor visibility. Thus, appropriate security policies should be

⁵¹ National Institute of Standards and Technology, "Security Strategies for Microservices-based Application Systems," (2019), <https://csrc.nist.gov/publications/detail/sp/800-204/final>.

⁵² National Institute of Standards and Technology, "Building Secure Microservices-based Applications Using Service-Mesh Architecture," (2020), <https://csrc.nist.gov/publications/detail/sp/800-204a/final>.

⁵³ National Institute of Standards and Technology, (2021), "Attribute-based Access Control for Microservices-based Applications using a Service Mesh," <https://csrc.nist.gov/publications/detail/sp/800-204b/final>.

⁵⁴ National Institute of Standards and Technology, "Implementation of DevSecOps for a Microservices-based Application with Service Mesh," (2022), <https://csrc.nist.gov/publications/detail/sp/800-204c/final>.

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 encryption in transit 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 to 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 microservices-based 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.⁵⁵ 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. See the Microservices scenario in Appendix A for additional details.

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.

⁵⁵ 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.

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 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 scale 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 in the case when an agency shifts services from a monolithic deployment. Agencies also should avoid oversimplification of microservices through excessive functional division, as this overcompensation could create more overhead and undermine the return-on-investment of this paradigm. At the same time, microservices will introduce operational complexity as many independent services will support a single application; so, agencies should also understand the level of visibility available to their monitoring services with respect to their microservices configurations.

Cloud-Native Authentication and Authorization

A common approach to deploying a microservices architecture and managing complexity is to use a service mesh. A service mesh is a dedicated infrastructure layer that provides configuration of network policies, traffic flow, and communication between services independent of the application code. Service meshes are delivered through “sidecar proxies” that are deployed on a per-application basis and run concurrently to each application. Generally, these sidecars are implemented as containers that are distinct from the microservices applications themselves. The service mesh must provide a means by which traffic may be routed to and from the application, in the form of ingress and egress gateways or via the sidecar proxies themselves. Security policies can then be enforced through these sidecars during runtime. The service mesh must act as a certificate authority (CA) for its sidecars and support an X.509 certificate infrastructure; this design ensures that all service traffic is encrypted in transit, a crucial security operational measure. However, an agency should not use a self-signed certificate to encrypt traffic in a production cloud environment, despite the presence of such a capability in certain service mesh implementations.

Cloud-native deployments that leverage a service mesh may use an attribute-based access control (ABAC) framework. ABAC is implemented through multiple functional modules—organized into well-defined architectures—that define and enforce access controls between a user and a protected resource. At their core, these modules define attributes based on user-object relationships and restrict how a user may interact with that object.

ABAC, in conjunction with a service mesh, allows for the definition and enforcement of granular and robust authorization and authentication frameworks for data moving within applications. The service mesh should have a decoupled control plane that encodes and distributes defined security policies for the microservices architecture. An agency may then support authorization and authentication through the control plane.

Authorization policies can be defined at the service or end-user level or based on access control models in the control plane of the service mesh. These policies are then pushed to the sidecar proxies that enforce them. These policies specify conditions upon which access may be allowed or blocked based on request metadata. Examples include source- or destination-based metadata such as internet protocol (IP) addresses or ports or hypertext transfer protocol (HTTP) request parameters or attributes. Additionally, the authorization framework must support the three principles comprising a reference monitor concept:

1. Invocation of the authorization mechanism on every access attempt (provided by the ingress/egress gateways and sidecar proxies);
2. Modification protection (provided by the ABAC modules that are separated from application logic and immutable); and
3. Correctness (through independent testing and verification of each module in both shadow IT and production).

In turn, authentication may be supported at the service or end-user levels. Service-based authentication is performed through service identity profiles, whereas end-user authentication is provided through supply of credentials. End-user authentication must be enforced by the sidecar proxy in a service mesh.

Agencies should look to evaluate access control solutions in terms of performance, flexibility, extensibility, scalability, and process isolation; and agencies should consider which software stack is most appropriate at each layer for their specific purposes. By combining ABAC policies with service meshes, agencies can more effectively manage their microservices architectures and authentication and authorization needs as their users and resources scale in the cloud.

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).

Agencies can use **Continuous Monitoring and Alerting** capabilities to validate their log usage and gain insight into their log statistics to ensure they are logging necessary data. Additionally, agencies may leverage these monitoring capabilities to ensure incoming log volume does not overwhelm log ingestion resources, as well as to create custom triggers on anomalous events. Agencies can utilize cloud **AI/ML** capabilities to filter log and telemetry data by removing noise and to identify anomalous traffic based on behaviors and historical data. Cloud-provided AI tools can train based on information from the CSP, such as traffic patterns and threat discovery, to improve logging functionality and adapt response procedures to changes in telemetry for agencies (that would otherwise be unachievable). **DevSecOps Integration** capabilities enable agencies to capture logs from pre-deployment in the CI/CD pipeline through end-user service in CDNs, increasing the scope of telemetry and logs well beyond the typical service time.

When collecting logs from SaaS, PaaS, or IaaS cloud instances, agencies should comply with the logging requirements issued by OMB M-21-31 pursuant to Section 8 of Executive Order 14028. 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:

- Ensure identity services are properly monitored for anomalous authentication and login attempts, especially around “break glass” accounts, privileged management/role changes, and key/secret vault changes.
- Monitor access policies and alert rules for undesired changes and monitor API activity logs and service metrics for anomalous activity.
- Perform routine system management, including data loss prevention, log maintenance, and monitoring for unexpected changes to logging policy.
- Detect cloud environmental changes in production applications, data/log storage, and cloud network through detection and prevention services, access managers, firewall, web application firewall, flow, and DNS records.

Time Synchronization

Agencies should ensure all collected logs meet minimum requirements and are recorded in 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 and can further amplify existing security concerns associated with latency. 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 use a default time zone for all log timestamps. 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.

On-Premises via Cloud Logs/Telemetry/Forensics

Many differences may exist between data collected on-premises and data collected via the cloud. Log delivery from CSPs generally have latencies that may reach 15 minutes or more before being made available. CSPs may not provide all the telemetry agencies had available for their on-premises operations. Agencies may not in some cases and will not in other cases have access to forensic artifacts, such as memory snapshots of machines suspected of being compromised. Agencies must be aware of these types

of differences and their impacts on their current security operations center (SOC), threat hunting, and incident response processes and procedures.

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, the data structures of the logs they use, mandate rate limits to prevent DoS attacks, and monitor API activity for future measurements and reporting. Agencies may want to consider designing webhooks to help reduce the load on API calls for event-based infrastructure.

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. 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 and PaaS

In IaaS and PaaS deployments, 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 and PaaS 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 work with developers to create tests that run as part of unit integration and/or system tests to certify the new deployments meet security standards. The security personnel of the DevSecOps team also work to ensure automated tests are in place to assess for common application vulnerabilities prior to

deployment. Security personnel work in collaboration with developers during the design process to ensure appropriate security practices are applied, and they also work with operations personnel to ensure the deployment is secure, properly monitored, and patched in a timely manner. Throughout the cyclical DevSecOps process, security personnel monitor for security issues. See Section 4.4 for additional details on DevSecOps.

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 (i.e., deploying upgrades without halting current operations). Adaptive AI/ML capabilities combined with proper ICAM capabilities may lead agencies to improved response times and a higher degree of fidelity in the deployment and orchestration of IaC.

To accomplish this, agencies can create base or “golden” VM 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 fulfills all appropriate security considerations. 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, hourly, or even in response to recently discovered vulnerabilities. For example, CI/CD pipelines should also address the use of vulnerable configurations, packages, and libraries within codebases by taking steps to alert and remediate. 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). In this deployment process, agencies should be aware of supply chain concerns with open source tools and should use a vetting solution to ensure library dependency versions are "secure/updated."

The cloud also allows agencies to delegate many maintenance tasks to the CSP, who offers IaaS, PaaS, and SaaS 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.

Agencies should develop, configure, and deploy with an IaC mindset. IaC allows agencies to manage and deploy configuration settings for everything from CSP-managed services to VMs and networks. Many CSPs offer tools to script and manage IaC, and there is third-party vendor software which may work across multiple clouds. Agencies should adopt best practices for using configuration management tools to

store and manage code including IaC code.⁵⁶ For example, repositories for code should not include sensitive information such as keys, emails, and passwords. Version control systems are one way to manage configurations asynchronously across systems.

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 may seek to use an alternative 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. Agencies should consider implementing separation of duties to ensure that no individual has access to encrypted content, keys, policies, and monitoring simultaneously. 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. Agencies should also determine how secrets will be stored either in a hardware (e.g., hardware security module (HSM)) or software setting (e.g., time-based one-time password (TOTP) authenticator application) and weight options in accordance with their threat model.

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. Please see Section 4.4 for more information on configuration management.

6. Conclusion

This Cloud Security Technical Reference Architecture illustrates recommended approaches to cloud migration and data protection for federal agencies as they continue to adopt cloud technology. These approaches will allow the Federal Government to identify, detect, protect, respond, and recover from cyber incidents, while improving cybersecurity across the .gov enterprise. Additionally, these approaches inform agencies on the advantages and inherent risks of adopting cloud-based services as their network architectures evolve.

The **Shared Services** section (Section 3) provided an overview of cloud service models and explained how agencies can leverage FedRAMP services to support their cloud migration. The **Cloud Migration** section (Section 4) highlighted various considerations for agencies as they design, implement, and maintain services in the cloud and included various scenarios to ensure efficient and secure migration to the cloud. Lastly, the **Cloud Security Posture Management** section (Section 5) introduced CSPM

⁵⁶ National Institute of Standards and Technology. "Draft NIST SP 800-204C Implementation of DevSecOps for a Microservices-based Application with Service Mesh," (2021), <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-204C-draft.pdf>.

capabilities, various cybersecurity outcomes that they support, and select applications to support agencies secure management of cloud resources, applications, and data, while also facilitating adoption of zero trust security principles.

This Cloud Security Technical Reference Architecture supports the continued evolution of federal agencies within a rapidly evolving technology landscape through a focus on cloud modernization.

Appendix A – Scenarios

The following three scenarios provide additional details associated with the adoption of federated identity management, microservices, and a warm standby site in the cloud. They are intentionally narrow in scope and are not intended to cover all possible implementations.

Federated Identity Management

Identity management is a critical component to enterprise security. As agencies move to the cloud, decisions must be made on how to manage identities across the many domains, services, and applications used.

Historically, software was purchased from vendors and installed in a traditional enterprise environment. Agencies are now moving beyond the traditional on-premises environment and consuming services from cloud service providers or from vendors operating their software as a service outside of an agency's environment. Without an integrated authentication solution, identity providers would be required for each distinct service environment causing each user in an agency to have multiple identities.

A solution to mitigate the burden of managing these multiple identities is federated identity management. By utilizing authentication standards such as the latest versions of SAML and OpenID, a single identity provider can be used across domains by applications and services as the source of authentication for identities. However, this single identity provider doesn't mean that an agency can – or should – use only one identity provider. Several factors should be considered to determine how many identity providers are needed to meet an agency's system requirements. The authentication standards establish a trust relationship between identity provider and each domain or service provider.

In Figure 15, a user requests access to a service. The service has a trust relationship with an identity provider that manages identities. Depending on how the authentication is implemented, the user may enter credentials at the service and the service will pass them to the identity provider, or the user may be redirected to the identity provider and then sent back to the service provider. The trust relationship between the identity provider and the service provider allows the service provider to accept the login by the user, whose credentials were verified by the identity provider.

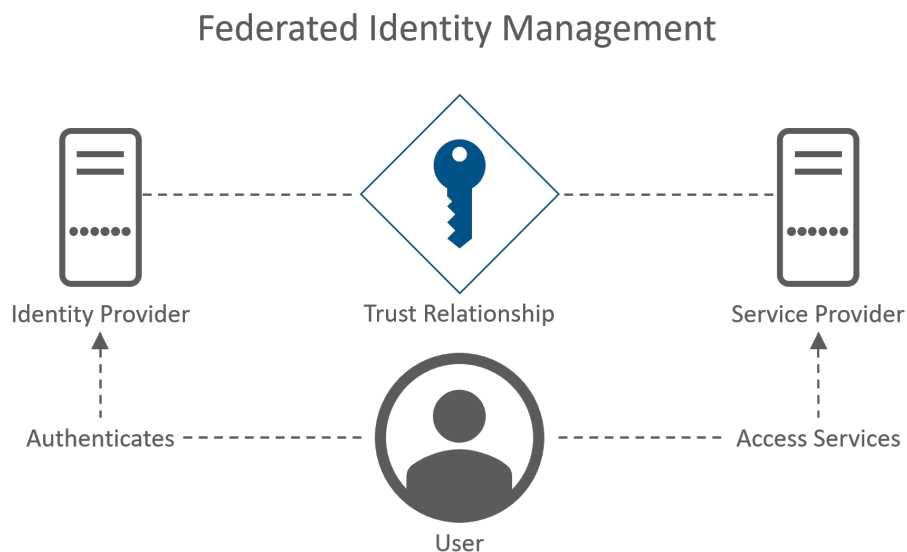


Figure 15: Federated Identity Management

Implementation Considerations:

- Single sign-on can also be implemented to reduce friction of workers who must pivot between applications and services in their jobs.
- Phishing-resistant multi-factor authentication can be integrated into federated identity management solutions.
- Identity management happens at one centralized location for the enterprise and not within domains or applications. This eases the management of identities when onboarding new personnel or shutting off access for departing personnel.
- Centralized authentication logs allow for rapid analysis of user activity to find suspicious logins or login attempts as opposed to needing services in place to collect authentication logs across domains or services for analysis.
- If compromised, threat actors can exploit the conveniences of a federated identity management systems, such as by exploiting a user's compromised credential to access other services.
- Not every service or application needs to utilize federated identity management. There may be circumstances where it is ideal to have a separate authentication realm for high security services, applications, or information.

Microservices

A well-established agency with mature development and DevOps teams wants to implement a zero-trust architecture (ZTA) as an integral part of its move to the cloud to both better secure their assets. The agency seeks to integrate this technology with its current infrastructure to minimize costs, but also to handle increasing demand for services and remain flexible to new requirements.

Traditionally, such an agency would leverage a monolithic architecture; any services added would have to modify a centralized codebase where, in general, changes are not easily scalable. Additionally, network policies would be rigidly defined with configuration manually performed on-premises. Configuration for services would likewise be performed on a per-device basis, which tends to introduce errors in consistency and policy management difficulties. By implementing a microservices architecture with complementary features such as a service mesh, configuration becomes centralized and can be pushed to networked devices uniformly, or granularly depending on agency requirements.

The agency decides to leverage a service mesh with a secure authentication and authorization framework to manage disparate services. Each microservice is untrusted, and so the service mesh with sidecar proxy provides additional security benefits that enhance independent development and deployment:

- In this scenario, the mesh provides the capability to deploy DevSecOps pipelines for IaC and policy-as-code, to incorporate security from the start. Microservices are atomic in nature and operate independently; each microservice performs a single, well-defined business function. As such, development of microservices is performed in a decentralized fashion, typically with small teams each contributing code for a service independently of other teams.
- The mesh complements a microservices-based architecture by compartmentalizing various cross-cutting stages of data analysis pipelines. This capability helps to address the collection and evaluation of vastly heterogeneous and unstructured data, and to do so at scale. The architecture can apply to different specialized domains, such as resource-constrained environments like IoT or SCADA.

Figure 16 depicts an example of such an implementation; the service mesh is implemented via sidecar proxies that are installed per-service (depicted via circles containing opposing arrows). Sidecar proxies

are applications that abstract certain features, such as inter-service communication, monitoring, and security, from the main architecture to make tracking and maintaining the application as a whole easier. This is the mechanism by which network and security policies may be pushed to microservices granularly. Each microservice may be developed by a distinct development team and uses its own dedicated data store. Agencies may access business functions through the API gateway, which manages interfaces for all microservices.

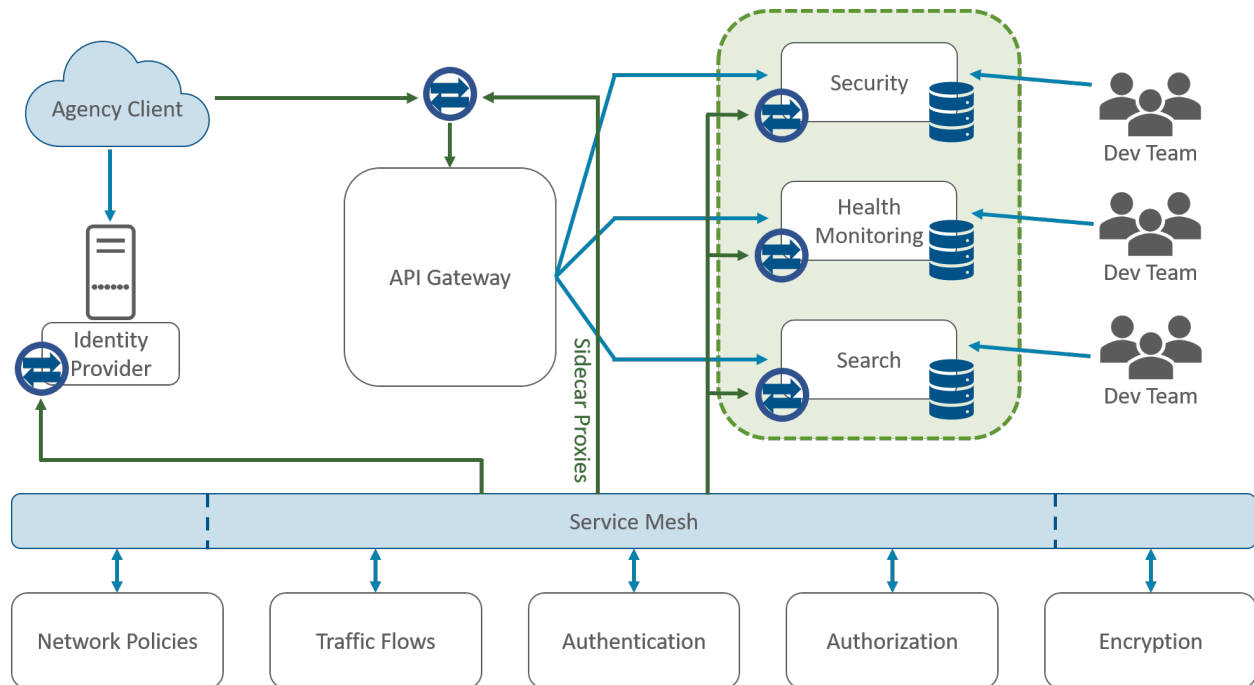


Figure 16: Microservices

Implementation Considerations:

- If an agency would like to integrate a security feature into a microservice, the agency should consider the requirements of each security service and understand the introduced risk. For example, introducing TLS encryption across containers in the above microservices-based architecture via a reverse proxy may introduce single points of failure. This should be weighed against the risks associated with unencrypted data in transit.
- Since the design of microservices inherently includes per-service independent development, data consistency may be an issue. Agencies should evaluate trade-offs between availability and consistency of the data and choose a strategy that is appropriate for their specific needs. This may include implementing a rolling data update strategy, whereby distributed data stores are evaluated and updated by a separate function for consistency.

Warm Standby

An agency would like to seamlessly transition workloads to a warm site in the cloud as needed during emergencies and high usage instances. This warm site must be updated regularly, and when possible after failovers, it must update the on-site live systems as well.

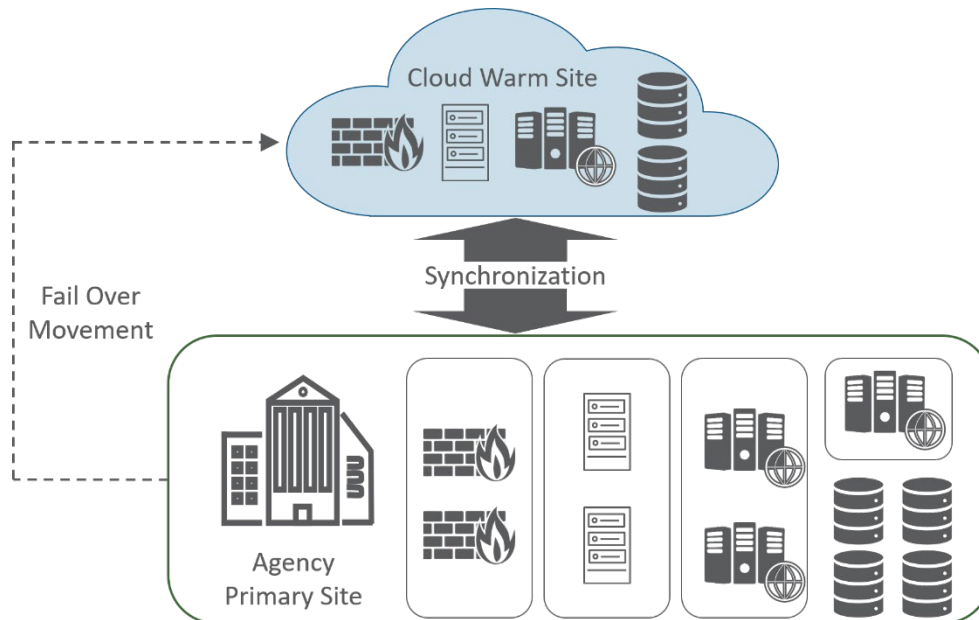


Figure 17: Cloud Warm Site Synchronization and Fail Over Movement

The cloud warm site should be synchronized to replicate security management, network access, service gateway, and data storage functions; however, the function of data manipulation should not occur in a warm site.

Typically, agencies seeking high availability from their operations will setup a hot site to complement performance and traffic which fail over from their primary systems. In these hot sites, replicas of infrastructure are synchronized immediately (notice the duplication of infrastructure in Agency Primary Site of Figure 17), and traffic is directed evenly to all replicas to increase network performance through load balancing. This load balancing also increases security through mitigating denial of service and implementing some moving target defense techniques.

As opposed to cold sites which include long-term storage with infrequent access and hot sites with full, immediate, mirrored tools, a warm site is implemented to enable continuity of operations during low availability or highly adversarial conditions. These warm sites operate at a reduced capacity from typical operations: only handling traffic and basic read-only requests while recovering systems and generating more replicas. Warm sites synchronize security management such as firewalls, network access such as routers, service gateways such as web servers, and service data with their primary counterparts at respectively decreasing regularity, for example, security management systems must be updated immediately to ensure proper configuration, whereas service data systems should not be immediately updated to prevent data corruption by adversaries from propagating.

Figure 17 highlights how a cloud-based warm site can be used to manage fail overs in addition to a traditional hot site fail over and load balancing system. Note that computation and manipulation of data should not be supported in the cloud warm. In this scenario IaaS is used for the warm site, though other service implementations could be deployed.

Implementation Considerations:

- A warm site ensures continuity of operations in worst-case scenarios while preserving original configurations and allowing the compromised environment may remain untouched, aiding in response and recovery.
- Data should be accessed in a primarily read-only state, both because writes may further corrupt sensitive data and because operating in disparate environments without real-time synchronization may lead to inconsistencies in data storage between cloud and traditional environments.
- By starting with implementing secondary fail over measures in cloud environments, agencies may leverage CSPM capabilities, such as Security and Risk Assessments and DevSecOps, to increase security without changes to existing infrastructure as the agency's network will be extended toward cloud-based warm sites.
- To ensure proper configuration and management prior to and during emergencies, an agency will need personnel familiar with synchronization, access management, capability implementation, and the general vulnerabilities and restrictions of CSP environments. These personnel will help agencies move further into cloud with future system implementations.

Appendix B – Glossary and Acronyms

This glossary contains cloud-specific terms and definitions that are used in this Technical Reference Architecture.

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.

Desktop-as-a-Service (DaaS): Desktop as a Service (DaaS) is a cloud computing offering where a service provider delivers virtual desktops to end users over the Internet, licensed with a per-user subscription.

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, Credential, and Access Management (ICAM): A fundamental and critical cybersecurity capability ensures the right people and things have the right access to the right resources at the right time for the right reason in support of federal business objectives.

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.

Supervisory Control and Data Acquisition (SCADA): A control system architecture comprising computers, networked data communications and graphical user interfaces for high-level supervision of machines and processes.

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

Service Mesh: A dedicated infrastructure layer that provides configuration of network policies, traffic flow, and communication between services independent of the application code. A service mesh supports a microservices architecture.

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.

Visibility: Refers to technical visibility (e.g., assets, users, systems, data, logs, etc.), operational visibility (e.g., usage, criticality, risks, etc.), and organizational visibility (e.g., mission functions, operations, priorities, etc.), and though one aspect may be specified, often a combination of the three are a concern.

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 C – 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.
- “Federal ICAM Architecture Introduction.” GSA. Accessed November 18, 2021.
<https://playbooks.idmanagement.gov/arch/>.
- 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/>.