



Trusted Internet Connections 3.0

Cloud Use Case

April 2023

Version 1.0

Cybersecurity and Infrastructure Security Agency

Cybersecurity Division

Disclaimer: This document is marked TLP:CLEAR. Disclosure is not limited. Sources may use TLP:CLEAR when information carries minimal or no foreseeable risk of misuse, in accordance with applicable rules and procedures for public release. Subject to standard copyright rules, TLP:CLEAR information may be distributed without restriction. For more information on the Traffic Light Protocol, see <http://www.cisa.gov/tlp/>.

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.

Revision History Table

Version	Date	Revision Description	Section/Page Affected
Draft	June 2022	Initial Release	All
1.0	April 2023	Request for Comments and Stakeholder Feedback	All

This use case references *Trusted Internet Connections 3.0 Security Capabilities Catalog*, v3.0, dated April 2023. The applicable security capabilities will be further explained in the document.

READER'S GUIDE

The Trusted Internet Connections (TIC) initiative is defined through key documents that describe the directive, the program, the capabilities, the implementation guidance, and capability mappings. Each document has an essential role in describing TIC and its implementation. The documents provide an understanding of how changes have led to the latest version of TIC and why those changes have occurred. The documents, which describe changes in architecture for TIC 3.0, are additive—each one builds on the one before—like chapters in a book. As depicted in Figure 1, the documents should be referenced in order and to completion to gain a full understanding of the modernized initiative.

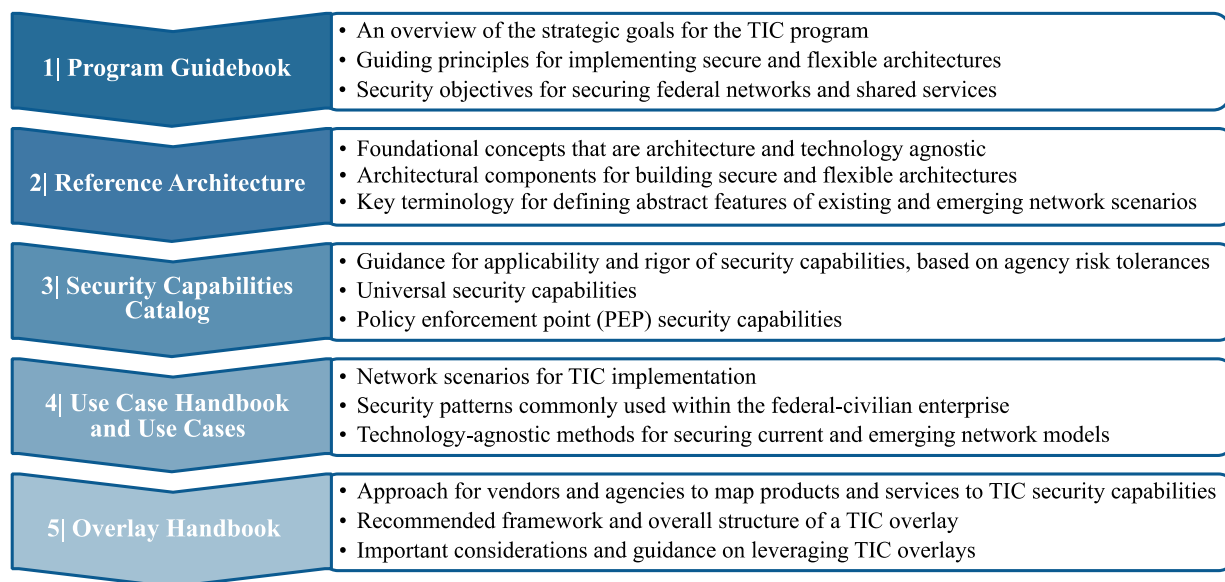


Figure 1 TIC 3.0 Guidance Snapshot

CONTENTS

1.	Introduction	1
1.1	Key Terms	1
2.	Overview of TIC Use Cases	2
3.	Purpose of the Cloud Use Case	3
4.	IaaS, PaaS, and SaaS Use Case	4
4.1	Assumptions and Constraints	4
4.2	Conceptual Architecture	7
4.3	Security Patterns	16
4.4	Applicable Security Capabilities	28
4.5	Telemetry Requirements	57
5.	Email-as-a-Service Use Case	58
5.1	Assumptions and Constraints	58
5.2	Conceptual Architecture	60
5.3	Security Patterns	64
5.4	Applicable Security Capabilities	70
5.5	Telemetry Requirements	90
6.	Conclusion	91
	Appendix A – Glossary and Definitions	92
	Appendix B – Related Federal Guidelines and Requirements	94
	Appendix C – Glossary for Cloud Use Case	95

Figures

Figure 1	TIC 3.0 Guidance Snapshot	iii
Figure 2	Use Case Trust Zone Legend	3
Figure 3	IaaS, PaaS, and SaaS Conceptual Architecture	7
Figure 4	Agency Cloud Services Communicating within a Single CSP and Region	8
Figure 5	Agency Cloud Services Communicating across Multiple Regions or Multiple CSPs	8
Figure 6	Varying Levels of Responsibilities for Different Service Models	11
Figure 7	Security Pattern 1 – Agency Campus to Cloud Service Provider	17
Figure 8	Security Pattern 2 – Remote User to Cloud Service Provider	18
Figure 9	Security Pattern 3 – External Entity to Cloud Service Provider	20
Figure 10	Security Pattern 4 – External Partners to Cloud Service Provider	22
Figure 11	Security Pattern 5 – Cloud Service Provider to the Web	24
Figure 12	Security Pattern 6 – Cloud Service Provider to Additional Agency Cloud Service	27
Figure 13	IaaS, PaaS, and SaaS Telemetry Sharing with CISA	57
Figure 14	EaaS Conceptual Architecture	61
Figure 15	Security Pattern 1 – Agency Campus User to Agency Email Service	65
Figure 16	Security Pattern 2 – Agency Remote User to Agency Email Service	67
Figure 17	Security Pattern 3 – External Entity to Agency Email Service	69
Figure 18	EaaS Telemetry Sharing with CISA	90

Tables

Table 1	Trust Zones in the Cloud Use Case for IaaS, PaaS, and SaaS	9
Table 2	Universal Security Capabilities for IaaS, PaaS, and SaaS	30
Table 3	Files PEP Security Capabilities for IaaS, PaaS, and SaaS	39
Table 4	Web PEP Security Capabilities for IaaS, PaaS, and SaaS	41
Table 5	Networking PEP Security Capabilities for IaaS, PaaS, and SaaS	44
Table 6	Resiliency PEP Security Capabilities for IaaS, PaaS, and SaaS	45
Table 7	Domain Name System PEP Security Capabilities for IaaS, PaaS, and SaaS	46

Table 8: Intrusion Detection PEP Security Capabilities for IaaS, PaaS, and SaaS	47
Table 9: Enterprise PEP Security Capabilities for IaaS, PaaS, and SaaS	49
Table 10: Data Protection PEP Security Capabilities for IaaS, PaaS, and SaaS	51
Table 11: Services PEP Security Capabilities for IaaS, PaaS, and SaaS	54
Table 12: Identity PEP Security Capabilities for IaaS, PaaS, and SaaS	55
Table 13: Trust Zones in the Cloud Use Case for EaaS	61
Table 14: Universal Security Capabilities for EaaS	72
Table 15: Files PEP Security Capabilities for EaaS	78
Table 16: Email PEP Security Capabilities for EaaS	79
Table 17: Domain Name System PEP Security Capabilities for EaaS	84
Table 18: Enterprise PEP Security Capabilities for IaaS, PaaS, and SaaS	84
Table 19: Data Protection PEP Security Capabilities for EaaS	85
Table 20: Identity PEP Security Capabilities for EaaS	87

1. INTRODUCTION

The initial versions of the TIC initiative sought to consolidate federal networks and standardize perimeter security for the federal enterprise. As outlined in *OMB Memorandum (M) 19-26: Update to the Trusted Internet Connections (TIC) Initiative*,¹ this modernized version of the initiative expands upon the original to drive security standards and leverage advances in technology as agencies adopt mobile and cloud environments. The goal of TIC 3.0 is to secure federal data, networks, and boundaries while providing visibility into agency traffic, including cloud communications.

1.1 KEY TERMS

To avoid confusion, terms frequently used throughout the TIC 3.0 documentation are defined below. Some of these terms are explained in greater detail throughout the TIC 3.0 guidance. A comprehensive glossary and acronyms list with applicable attributions can be found in Appendix A.

- **Boundary:** A notional concept that describes the perimeter of a zone (e.g., mobile device services, general support system (GSS), Software-as-a-Service (SaaS), agency, etc.), within a network architecture. The bounded area must have an information technology (IT) utility.
- **Internet:** The internet is discussed in two capacities throughout TIC documentation.
 1. A means of data and IT traffic transport.
 2. An environment used for web browsing purposes, hereafter referred to as “web.”
- **Managed Trusted Internet Protocol Services (MTIPS):** Services under GSA’s Enterprise Infrastructure Solutions (EIS) contract vehicle that provide TIC solutions to government clients as a managed security service. It is of note that the EIS contract is replacing the GSA Networkx contract vehicle that is set to expire in Fiscal Year (FY) 2023.
- **Management Entity (MGMT):** An entity that oversees and controls security capabilities. The entity can be an organization, network device, tool, service, or application. The entity can control the collection, processing, analysis, and display of information collected from the policy enforcement points (PEPs), and it allows IT professionals to control devices on the network.
- **Policy Enforcement Point (PEP):** A security device, tool, function, or application that enforces security policies through technical capabilities.
- **Security Capability:** A combination of mutually reinforcing security controls (i.e., safeguards and countermeasures) implemented by technical means (i.e., functionality in hardware, software, and firmware), physical means (i.e., physical devices and protective measures), and procedural means (i.e., procedures performed by individuals).² Security capabilities help to define protections for information being processed, stored, or transmitted by information systems.
- **Telemetry:** Artifacts derived from security capabilities that provide visibility into security posture.

¹ “Update to the Trusted Internet Connections (TIC) Initiative,” Office of Management and Budget M-19-26 (2019), <https://www.whitehouse.gov/wp-content/uploads/2019/09/M-19-26.pdf>.

² “SP 800-53 Security and Privacy Controls for Federal Information Systems and Organizations,” National Institute of Standards and Technology (December 2020), <https://csrc.nist.gov/publications/detail/sp/800-53/rev-5/final>.

- **TIC:** The term “TIC” is used throughout the Federal Government to denote different aspects of the TIC initiative; including the overall TIC program, a physical TIC access point (also known as a Traditional TIC), and a TIC Access Provider (TICAP – see below). This document refers to TIC as an adjective or as the Trusted Internet Connections initiative.
- **TIC Access Point:** The physical location where a federal civilian agency consolidates its external connections and has security controls in place to secure and monitor the connections.
- **TIC Access Provider (TICAP):** An agency or vendor that manages and hosts one or more TIC access points. Single Service TICAPs serve as a TIC Access Provider only to their own agency. Multi-Service TICAPs also provide TIC services to other agencies through a shared services model.
- **TIC Overlay:** A mapping of products and services to TIC security capabilities.
- **TIC Use Case:** Guidance on the secure implementation and/or configuration of specific platforms, services, and environments. A TIC use case contains a conceptual architecture, one or more security pattern options, security capability implementation guidance, and CISA telemetry guidance for a common agency computing scenario.
- **Trust Zone:** A discrete computing environment designated for information processing, storage, and/or transmission that share the rigor or robustness of the applicable security capabilities necessary to protect the traffic transiting in and out of a zone and/or the information within the zone.
- **Web:** An environment used for web browsing purposes. Also see **Internet**.

2. OVERVIEW OF TIC USE CASES

TIC use cases provide guidance on the secure implementation and configuration of specific platforms, services, and environments, and will be released on an individual basis. The guidance is derived from pilot programs and best practices from the public and private sectors. The purpose of each TIC use case is to identify the applicable security architectures, data flows, and PEPs, as well as describe the implementation of the security capabilities in a given scenario. TIC use cases articulate:

- Network scenarios for TIC implementation,
- Security patterns commonly used within the federal civilian enterprise, and
- Technology-agnostic methods for securing current and emerging network models.

TIC use cases build upon the key concepts and conceptual implementation of TIC 3.0 presented in the *TIC 3.0 Reference Architecture* (Reference Architecture) and provides implementation guidance for applicable security capabilities defined in the *TIC 3.0 Security Capabilities Catalog* (Security Capabilities Catalog). The *TIC 3.0 Use Case Handbook* (Use Case Handbook) provides general guidance for how agencies can utilize and combine use cases.

Agencies have flexibility in implementing TIC use cases. In particular:

- An agency may combine one or more use cases to best design and implement their TIC architectures.
- Use cases may provide more than one option for implementing a security pattern in order to give agencies flexibility.
- Each trust zone in a use case will be labeled with a high, medium, or low trust level based on a pilot implementation or best practice. The use cases are depicted following the schema illustrated in Figure 2. Agencies can modify this trust zone designation to meet their needs and reflect their environment, including assigning a zone to a different trust level or altering the number of trust levels and their labels. Refer to the Reference Architecture for more details on trust zones.



Figure 2: Use Case Trust Zone Legend

- When securing trust zones, agencies should consider unique data sensitivity criteria and the impact of compromise to agency data stored in trust zones. Agencies may apply additional security capabilities that have not been included in the use case.
- Agencies have the discretion to determine the level of rigor necessary for applying security capabilities in use cases, in accordance with federal guidelines and agency risk tolerance.

Refer to the Use Case Handbook for more information on TIC use cases.

3. PURPOSE OF THE CLOUD USE CASE

The *TIC 3.0 Cloud Use Case* (Cloud Use Case) defines how network and multi-boundary security should be applied in cloud environments. The use case is broken into two distinct parts, focusing on cloud deployments for:

1. Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS), and Software-as-a-Service (SaaS) (Section 4), and
2. Email-as-a-Service (EaaS) (Section 5).

Appendix C contains definitions of common terms that are used to describe cloud computing throughout this use case.

According to Executive Order 14028,³ “Improving the Nation’s Cybersecurity,” the Federal Government must adopt security best practices; advance toward Zero Trust Architecture; accelerate movement to secure cloud services, including Software as a Service (SaaS), Infrastructure as a Service (IaaS), and Platform as a Service (PaaS). Additionally, the OMB Zero Trust Strategy Memo (M-22-09)⁴ encourages agencies to use the security features in cloud infrastructure, requires agencies to meet certain cybersecurity baselines for zero trust, and have a long-term implementation plan in place to move towards a zero trust architecture. This use case can be used by agencies to make use of cloud infrastructure and to secure their SaaS, IaaS, PaaS, and EaaS environments. While this use case can be leveraged as agencies move towards zero trust architectures, implementation of zero trust requires additional controls, additional rigor of applying security capabilities, and measures beyond those detailed in this use case.

The recent *Cloud Security Technical Reference Architecture*⁵ provides strategic and technical guidance to agencies as they adopt cloud technology. This use case leverages the shared security model and cloud security

³ “Executive Order 14028 Improving the Nation’s Cybersecurity,” Office of Management and Budget (May 2021), <https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/>.

⁴ “Moving the U.S. Government Toward Zero Trust Cybersecurity Principles,” Office of Management and Budget M-22-09 (2022). <https://www.whitehouse.gov/wp-content/uploads/2022/01/M-22-09.pdf>.

⁵ “Cloud Security Technical Reference Architecture,” Cybersecurity and Infrastructure Security Agency, United States Digital Service, and General Services Administration (2021). <https://www.cisa.gov/cloud-security-technical-reference-architecture>.

posture management guidance in the *Cloud Security Technical Reference Architecture*, as applicable to applying TIC security capabilities.

Specifically, this use case is from the vantage point of cloud-hosted services. Information from the vantage point of the client accessing the cloud-hosted services can be found in other use cases, including the *TIC 3.0 Branch Office Use Case* (Branch Office Use Case) and *TIC 3.0 Remote User Use Case* (Remote User Use Case).

4. IAAS, PAAS, AND SAAS USE CASE

This section broadly covers IaaS, PaaS, and SaaS deployment models, as outlined in OMB M-19-26. This section does not detail specific SaaS applications, such as EaaS or Unified-Communications-as-a-Service, or specific PaaS services, such as Managed Container Orchestration Services or Function-as-a-Service (FaaS). This section includes six network security patterns:

- Secure agency campus to agency-sanctioned cloud service providers (CSPs)
- Secure remote user access to agency-sanctioned CSPs
- External entity accessing agency CSP services
- Secure agency CSP service accessing resources from external partners
- Secure agency CSP service accessing resources in the web
- Agency cloud service to agency cloud service

An agency may implement a subset of these security patterns and not necessarily all six, depending on how agencies are migrating and deploying services in the cloud. For example, an agency may not have agency CSP services accessible by external entities.

Agencies may implement **additional security patterns** not covered in the Cloud Use Case.

Agencies may implement additional security patterns. These additional security patterns may be in scope for a different use case but would be out of scope with respect to the IaaS, PaaS, and SaaS guidance in the Cloud Use Case.

4.1 ASSUMPTIONS AND CONSTRAINTS

This section outlines guiding assumptions and constraints for the IaaS, PaaS, and SaaS guidance in the Cloud Use Case. It is intended to clarify significant details about the construction and replication of the IaaS, PaaS, and SaaS guidance in this use case. The assumptions are broken down by IaaS, PaaS, and SaaS guidance in this use case as a whole and by the unique entities discussed in this section:

- | | |
|--------------------------|---------------------|
| • Agency campus | • External partners |
| • Agency-sanctioned CSPs | • External entities |
| • Remote users | • Web |

The following are the assumptions and constraints of the IaaS, PaaS, and SaaS guidance in this use case:

- Requirements for information sharing with CISA in support of National Cyber Protection System (NCPS) and Continuous Diagnostics and Mitigation (CDM) purposes are beyond the scope of this document. Consult the NCPS program⁶ and CDM program⁷ for further details.
- Requirements for endpoint protection are beyond the scope of this document. Consult the Federal Information Security Modernization Act of 2014 (FISMA) or National Institute of Standards and Technology (NIST) references in Appendix B for additional guidance on endpoint protections, Bring Your Own Device (BYOD), and telework security.
- The TIC security capabilities applicable to the use case do not depend on a particular data transfer mechanism. In other words, the same capabilities apply if the conveyance is over leased lines, software virtual private network (VPN), hardware VPN, etc.
- The scope of the IaaS, PaaS, and SaaS guidance in the Cloud Use Case is primarily focused on network security. While this use case can be compatible with zero trust, implementation of zero trust requires additional controls and measures beyond those detailed in this use case.

The following are assumptions about the agency campus:

- For this use case, the agency campus entity may refer to the main agency campus, as well as branch offices.
- The agency campus utilizes the *TIC 3.0 Traditional TIC Use Case* (Traditional TIC Use Case), or equivalent security architectures, to access the web and CSPs.
- Any branch offices utilize the Branch Office Use Case, or equivalent security architectures, to access the web, CSPs, and the agency campus.
- The agency maintains control over, and has significant visibility into, the agency campus.
- Data is protected at a level commensurate with the agency's risk tolerance and in accordance with applicable federal requirements.
- The agency employs network operation center (NOC) and security operation center (SOC) tools capable of maintaining and protecting their portions of the overall infrastructure. To accomplish this, agencies can opt to use an NOC and SOC, or commensurate solutions.

The following are assumptions about agency-sanctioned CSPs:⁸

- CSPs are compliant with the applicable Federal Risk and Authorization Management Program (FedRAMP) requirements.⁹
- Interactions with CSPs follow agency-defined policies and procedures for business need justification, partner connection eligibility, service levels, data protections, incident response information sharing and reporting, costs, data ownership, Authority to Operate (ATO), and contracting.
- The agency maintains awareness of which CSPs and CSP services are sanctioned for use by the agency.
- The agency has limited control over and visibility into many CSP environments.
- CSPs have NOCs and SOCs that control and protect the portions of the service infrastructure where the agency has little or no control or visibility.
- The agency only uses secure mechanisms (e.g., transport layer security (TLS) or VPN) for CSP service administration.
- The agency only uses strong authentication mechanisms (e.g., Federal Information Processing Standard [FIPS] 140-3¹⁰ compliant multi-factor authentication (MFA) for CSP service administration.

⁶ "National Cybersecurity Protection System," Cybersecurity and Infrastructure Security Agency. <https://cisa.gov/national-cybersecurity-protection-system-ncps>.

⁷ "Continuous Diagnostics and Mitigation," Cybersecurity and Infrastructure Security Agency. <https://cisa.gov/cdm>.

⁸ See section 4.2.1 for further detail on the division of responsibility between Agencies and CSPs.

⁹ "FedRAMP," General Services Administration (2019). <https://www.fedramp.gov/federal-agencies/>.

¹⁰ "FIPS 140-3 NIST Security Requirements for Cryptographic Modules," National Institute of Standards and Technology (2019), <https://csrc.nist.gov/publications/detail/fips/140/3/final>.

- Data stored at CSPs are at a level commensurate with the agency's risk tolerance and in accordance with applicable federal requirements.
- CSPs allow the agency to define and/or configure policies that the CSP applies on their behalf.
- CSPs allow the agency to define roles and responsibilities for the definition and configuration of policies applied on their behalf by the CSP.
- CSPs provide the agency with mechanisms for obtaining visibility into the current state and history of the system (e.g., log information, configuration, accesses, system activity).
- CSPs provide commensurate protections and policy enforcement for traffic between the agency tenant and other tenants of the CSP as between the agency tenant and parties outside the CSP.

The following are assumptions about remote users:

- The remote user utilizes the Remote User Use Case, or equivalent security architectures, to access the agency campus, the web, and CSPs.
- The remote user may be using either government furnished equipment (GFE) or BYOD.
- For GFE, remote users may be permitted business only use of their devices (e.g., Corporate-Owned Business Only [COBE]), or permitted for personal use (e.g., Corporate-Owned Personally Enabled [COPE]).
- Devices employed by remote users may include desktops, laptops, and mobile devices (e.g., smartphones and tablets). While remote users may connect to virtual desktop instances hosted by the agency or in cloud service providers, these agency-managed desktop instances are not considered remote user devices. However, they may be considered as agency virtual GFEs inside an agency campus or cloud environment.
- For GFE, the agency maintains control over and has significant visibility into devices used by the remote user. All traffic from GFE devices is in scope for TIC 3.0.
- Guidance on BYOD policies is beyond the scope of this document; however, the agency may have limited control and visibility into BYOD devices. Traffic from BYOD used for official purposes to the agency campus and to agency-sanctioned CSPs is in scope for TIC 3.0. While traffic to the web from BYOD is generally out of scope for TIC 3.0, if traffic to the web originates from an application accessing agency data, then the traffic would be in scope for TIC 3.0.
- Traditionally, the remote user would have used the agency campus for all CSP and web traffic.
- Agency data on remote user devices, or in transit to and from them, is protected at a level commensurate with the agency's risk tolerance and in accordance with applicable federal requirements.
- The agency employs NOC and SOC tools capable of protecting remote user sessions. These functions may be performed as an extension to the NOC and SOC tools managed and housed at the agency campus or via commensurate solutions.

The following are assumptions about external partners:

- The agency's interactions with external partners follow agency-defined policies and procedures for business need justification, partner connection eligibility, service levels, data protections, incident response information sharing and reporting, costs, data ownership, and contracting.
- The agency uses only limited and well-defined services of external partners or permits external partners access to only limited and well-defined services of the agency.
- The agency has limited control over and visibility into the internal networks or information technology operations of external partners.
- External partners have NOCs and SOC that control and protect the portions of their infrastructure where the agency has little to no control or visibility.
- The agency only uses secure mechanisms (e.g., TLS) to communicate with external partners.

- The agency only uses strong authentication mechanisms (e.g., FIPS 140-3 compliant MFA) with external partners.¹¹
- Data provided to external partners is protected at a level commensurate with the agency's risk tolerance and in accordance with federal requirements.

The following are assumptions about external entities:

- External entities include public users accessing agency services.
- The agency may not be able to rely on policies deployed by external entities.

The following are assumptions about the web:

- The web contains untrusted entities.
- The agency has no ability to apply policy in the web or to web resources.

4.2 CONCEPTUAL ARCHITECTURE

The IaaS, PaaS, and SaaS guidance in the Cloud Use Case focuses on the scenario in which an agency has one or more cloud deployments in its enterprise. Traditionally, agency users have accessed cloud deployments either directly from an agency campus or by establishing a secure connection (e.g., VPN) to an agency campus and using that channel to access the cloud deployment.

As shown in Figure 3, this conceptual architecture is composed of eight distinct trust zones: agency campus, CSP, remote user, external partner, agency service, external entity, web, and an additional agency cloud service. This conceptual architecture also shows a single remote user, a single external partner, and a single external entity trust zone. These simplifications are not meant to imply that an agency must treat all remote users, external partners, external entities, or CSPs in the same manner. Applicable TIC capabilities and their rigor should be tailored for the nature of the remote user, external entity, or the CSP service in use.

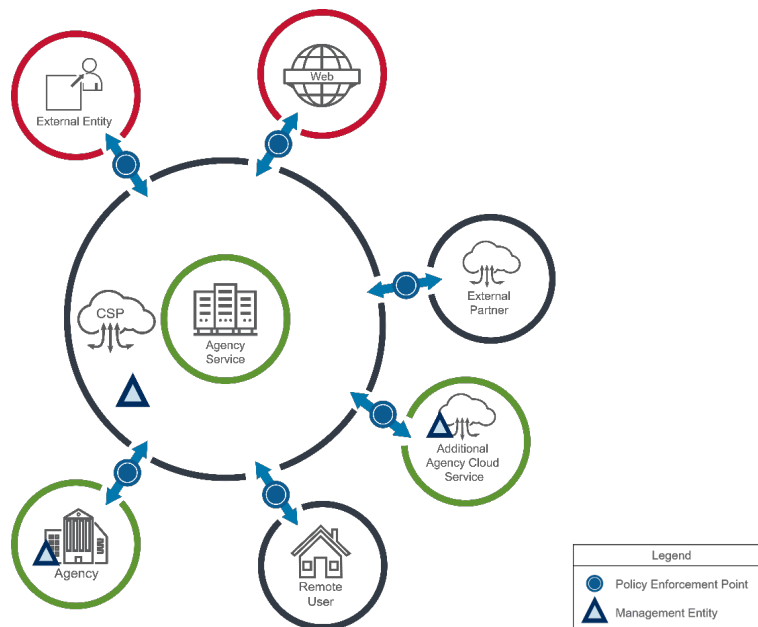


Figure 3: IaaS, PaaS, and SaaS Conceptual Architecture

¹¹ "FIPS 140-3 NIST Security Requirements for Cryptographic Modules," National Institute of Standards and Technology (2019). <https://csrc.nist.gov/publications/detail/fips/140/3/final>.

The conceptual architecture in Figure 3 shows two separate agency service trust zones: an agency service and an additional agency service. This is done to facilitate the discussion of the security pattern in Section 4.3.6, where two agency cloud services are communicating. The deployment of these trust zones will vary depending on agency design decisions and architecture considerations. The following two detailed figures depict different scenarios that agencies may consider.

In the scenario depicted in Figure 4, the two agency services may be communicating across logical boundaries within a single tenant or between tenants. Traffic between the two services may leave the Agency's trust boundary, but generally does not leave the CSP's private network. Security controls will be broadly similar on both sides, but may vary if the service model (IaaS, PaaS, SaaS) differs.

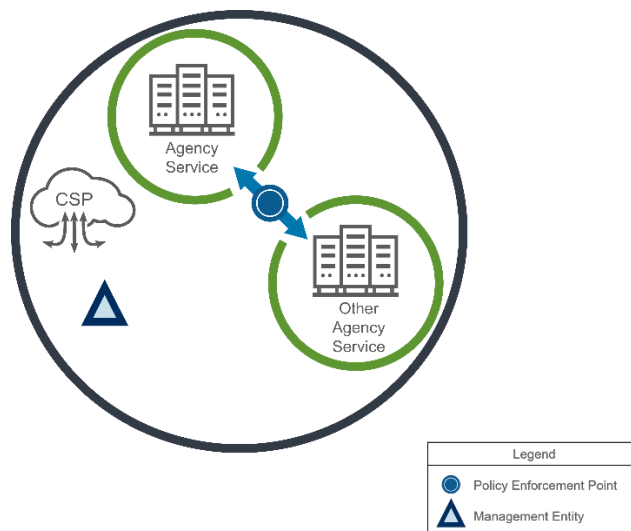


Figure 4: Agency Cloud Services Communicating within a Single CSP and Region

In the scenario depicted in Figure 5 the two agency services are communicating across a regional or CSP boundary. Traffic between the two services will leave both the Agency and CSP trust boundaries. Security controls may differ substantially if the additional service is hosted in a different CSP.




Figure 5: Agency Cloud Services Communicating across Multiple Regions or Multiple CSPs

The trust zones depicted in Figure 3 are detailed in Table 1. The trust zones are labeled with levels of trust, using the example trust levels—high, medium, and low—explained in the Reference Architecture. While the trust levels assigned to each of these zones in Table 1 were selected based on existing pilots or deployments, the trust assignments may not capture the needs or requirements of all agencies. Agencies may assign different trust levels to trust zones, based on their own risk tolerance. For example, an agency might decide to designate a CSP with a higher trust level based on agency criteria (e.g., the accreditation level of the CSP, the control and visibility, available protections). Additionally, an agency may have remote users that employ unmanaged personal devices and may decide to label remote users with a lower trust level.

Implementation Consideration

The trust levels in this use case are intended to be examples. Agencies may define and assign trust levels to align with their requirements, environments, and risk tolerance.

Table 1 briefly explains why each entity is labeled with either a high, medium, or low trust zone level in this conceptual architecture to help agencies determine what is most appropriate in their implementation.

Table 1: Trust Zones in the Cloud Use Case for IaaS, PaaS, and SaaS

Trust Zone	Description
Agency Campus Trust Zone	The Agency Campus Trust Zone is the logical zone for the agency campus or the agency's enterprise network. The trust zone includes management entities (MGMTs) such as the NOC, SOC, and other entities. The agency maintains control over and visibility into the agency campus. The agency campus employs the Traditional TIC or Branch Office Use Cases, or equivalent, including when transmitting traffic from the CSP to external entities. The Agency Campus Trust Zone is labeled with a high trust level in this conceptual architecture.
Cloud Service Provider Trust Zone	The Cloud Service Provider Trust Zone is a logical trust zone for the CSP providing IaaS, PaaS, SaaS, or a similar service. The CSP environment follows a shared responsibility model, with the CSP responsible for protecting the underlying cloud infrastructure and the agency providing certain policy-defined functions and capabilities. The trust zone includes a MGMT that executes locally scoped functions for the CSP environment. The Cloud Service Provider Trust Zone is labeled with a medium trust level in this conceptual architecture due to the potential for limited agency control over and visibility into the CSP environment.
Remote User Trust Zone	The Remote User Trust Zone is a logical trust zone representing a device employed by a remote user when accessing agency resources. Remote user devices may be agency-managed (e.g., GFE) or not managed by agencies (e.g., BYOD). Devices not managed by agencies may not be suitable for performing some policy enforcement capabilities. The agency may have no control over or visibility into non-GFE devices and may have limited control over or visibility into agency-managed devices. The remote user employs the Remote User Use Case. The Remote User Trust Zone is labeled with a medium trust level in this conceptual architecture.
External Partner Trust Zone	The External Partner Trust Zone is a logical trust zone for an external partner that offers services to or receives services from the agency. The agency has limited control over and visibility into the external partner environment. The agency can provide certain defined capabilities for an external partner to manage, and the external partner is responsible for protecting the underlying infrastructure. The trust zone may include a MGMT with functions locally scoped for the environment. Policy enforcement between the external partner and the CSP or between the external partner and the agency campus may use a shared responsibility model. Given the more limited control and visibility available to the agency, the External Partner Trust Zone is labeled with a medium trust level in this conceptual architecture.

Trust Zone	Description
Agency Service Trust Zone	The Agency Service Trust Zone is a logical trust zone that represents a service that an agency deploys in the cloud. This may be an IaaS, PaaS, SaaS, or similar service. The agency service has a shared responsibility model, with the CSP responsible for protecting the underlying cloud infrastructure and the agency providing policy-defined functions and protections in line with the agency risk tolerances. The Agency Service Trust Zone is labeled with a <i>high trust level</i> in this conceptual architecture because the agency has control over and visibility into the service.
Additional Agency Service Trust Zone	The Additional Agency Service Trust Zone is a logical trust zone representing a service, different from the Agency Service Trust Zone service, that an agency deployed into a different cloud environment or potentially into a different tenancy or region in the same cloud environment. This may be an IaaS, PaaS, SaaS, or similar service, and may differ from the deployment model used for the service in the Agency Service Trust Zone. The Additional Agency Service Trust Zone is labeled with a <i>high trust level</i> in this conceptual architecture because the agency has control over and visibility into the service.
External Entity Trust Zone	The External Entity Trust Zone is a logical zone that depicts an untrusted and unmanaged user of agency services with no PEPs or MGMTs where the agency, or entities acting on its behalf, may deploy policies. An external entity may also depict a nonhuman entity. Given these limitations, the External Entity Trust Zone is labeled with a <i>low trust level</i> in this conceptual architecture.
Web Trust Zone	The Web Trust Zone is a logical trust zone that depicts an environment with untrusted external resources, including non-agency-sanctioned cloud service providers that agency users or entities may access from or through the CSP. This trust zone does not permit the agency nor entities acting on its behalf to deploy or enforce policies. Given these limitations, the Web Trust Zone is labeled with a <i>low trust level</i> in this conceptual architecture.

4.2.1 Shared Security Model

This use case provides capability guidance for SaaS, PaaS, and IaaS. Each of these service offerings has differences in how security protections are managed. This is commonly represented via a shared security model, as illustrated in Figure 6. An agency needs to understand this model and what security protections are handled by each CSP versus the agency to fulfill both universal and PEP security capabilities and to ensure parity across all PEPs.

The responsibility for securing a SaaS offering relies heavily upon the service provider. On the other hand, with IaaS, most responsibility falls on the agency, some responsibility resides with the CSP, and other responsibilities are shared. While the shared responsibility shows three distinct service models, as cloud offerings mature, there is no clean line between offerings and the delineation between each service model is blurred. Additionally, each CSP may define this shared security relationship differently. Agencies must clearly identify and understand the delineation of responsibilities between themselves and their CSPs for deploying security capabilities. This can become more complex when agencies are utilizing services from multiple CSPs.



Figure 6: Varying Levels of Responsibilities for Different Service Models

4.2.2 Risk and Deployment Considerations

As agencies migrate data and applications from on-premises deployments to cloud deployments, they must understand the differences between the two models, how to protect new cloud deployments, how the agency security posture must adapt, and best practices for mitigating any risks.¹²

4.2.2.1 Control and Visibility

With agency data and services hosted on-premises, agency workflows and processes often assume physical access to, and control of, all hardware, networks, and facilities. When agencies move data and services to the cloud, they will no longer have physical access, and may have potentially different levels of control and visibility. Therefore, agencies should understand any potential differences, as well as how agency workflows and processes can be augmented or changed to account for them. Agencies should understand the controls in place for physical security at a CSP and who has access to physical devices containing agency data. Additionally, agencies should understand the security policies of CSPs and their alignment with agency risk tolerances and, if needed, consider service level agreement (SLA) language or other security protections to better align.

When agencies move from on-premises to the cloud, they may have less visibility into the CSP supply chain for hardware or other services. In PaaS or IaaS deployments, agencies may be able to maintain awareness of software versioning (e.g., operating systems, web servers). However, in a SaaS deployment, the agency often does not have visibility into the hardware and software that supports the services. Agencies should ensure that their supply chain risk management strategies are in alignment with such differences in visibility.

In traditional deployments, agency policies and procedures for maintaining the confidentiality, integrity, availability and recoverability of agency data, and data backups often had dependencies on the agency's ability to control the location where data and backups are stored, when data backups are performed, what recovery plans are in place in case of a data breach or accidental loss, and what happens to deleted data. As agencies

¹² National Institute of Standards and Technology Cybersecurity Framework, <https://www.nist.gov/cyberframework/nist-cybersecurity-framework-csf-reference-tool>, and National Institute of Standards and Technology Risk Management Framework, <https://csrc.nist.gov/Projects/risk-management/about-rmf>.

migrate data into cloud environments, they will need to work with the CSPs to understand and, if needed, augment the data security measures to ensure alignment with agency risk tolerance, applicable federal requirements, and to update their policies and procedures to account for the differences between the cloud and their on-premises environments.

Traditional agency incident response procedures are often aligned with the types of visibility and control in traditional deployments. Agencies will need to understand potential differences in visibility and control that they have in cloud environments, and how to adapt their incident response procedures to account for such. Additionally, there may be opportunities to increase visibility in the cloud environment, potentially by employing different tiers of service, or supplementing visibility by deploying additional security protections or monitoring solutions, either native CSP and/or third-party tools. Agencies should understand the visibility capabilities available for CSPs and how best to leverage them commensurate with the risks to agency services or data.

4.2.2.2 Service and Cloud Availability

In a traditional deployment, services and data are deployed on agency-controlled infrastructure in agency environments, often co-located with their agency campuses. Agencies often used this positioning as part of ensuring the availability of services and data aligned with the needs of their on-premises users. While the transition to cloud-hosted services can potentially improve that availability for both on-premises and remote users (e.g., by providing more robust availability zones and a more diverse physical infrastructure footprint), agencies may need to ensure their procedures and processes account for potential differences in maintaining availability, technical issues (e.g., the loss of connectivity to the CSP), and issues that can potentially require agencies to change CSPs (e.g., CSP going out of business or changing business models).

4.2.2.3 Use of Shared Infrastructure

Agencies often used the control provided in traditional deployment models to control the access to and use of the infrastructure their services support. With the transition to a cloud environment, the infrastructure used to deploy or support agency services is controlled by a third party and is often shared with other organizations. The security of this shared infrastructure, thus, depends on effective separation of tenant deployments so that malicious entities affecting one tenant cannot gain access to, or otherwise influence, other tenants. Cloud providers often provide options for letting tenants better define their isolation (e.g., hardware-level isolation, single tenant infrastructure, SLAs), and agencies will need to understand the options available from the provider as they align their cloud deployments according to their business needs and risk tolerances.

4.2.2.4 Identity and Access Control

In traditional enterprise deployments, agencies often centralized their user identities, with users having a single identity throughout the enterprise environment and developed user lifecycle procedures and access controls around this centralized identity. As agencies transition to cloud services, they will need to understand the opportunities for maintaining these centralized user identities and any updates that might be needed for their user lifecycle procedures.

Various methods for integrating cloud identities with the traditional enterprise identity, include single sign-on, federated identity, or the use of Identity-as-a-Service providers. This integration can help provide enterprise-wide visibility and consistency, though agencies will need to understand any potential this might open for lateral movement, enabling the compromise of one agency environment to be extended to other environments.

As agencies transition their identity to integrate cloud identities, they should work with their users during the transition as threat actors may attempt to exploit user confusion around these changes in identity to facilitate phishing and other attacks.

4.2.2.5 Management Plane Accessibility

Moving to the cloud carries a significant change in the way that agencies manage resources. While traditional management activities often involved specific devices behind layers of protection, cloud management interfaces are commonly application program interfaces (APIs) made available directly over the internet which can provide adversaries with additional opportunities to attempt to compromise or disrupt the management plane. Agencies will need to understand the protection opportunities available to ensure they can maintain comparable or improved security for their management process as they transition into these environments.

Management networks may also behave in unique ways internal to the cloud. In some cases, management network access may be implicit, bypassing security boundaries such as firewalls and meaning that resources may be less isolated than is immediately obvious. Agencies will need to understand the available methods for accessing cloud resources through the management plane and, if needed, augment their security protections to account for these mechanisms.

The methods of data exfiltration available in cloud environments may differ from those in traditional deployment models. Cloud environments often facilitate data transfer between cloud tenants in ways that may bypass agency data loss prevention policies based around traditional network-based models of exfiltration. For example, a threat actor may be able to directly share a backup of an agency environment with an attacker-controlled tenant or may be able to create a trusted relationship with an attacker-controlled tenant, allowing them to directly access any of the agency data. Agencies should understand the controls and architectures available for their CSPs to protect against data exfiltration and limit unauthorized access to the management plane. Agencies should ensure that their data loss prevention techniques can detect and potentially prevent attempts at exfiltration using these other techniques.

4.2.2.6 Misconfiguration

Agencies often make use of network architectures to decrease opportunities for misconfigured services and infrastructure to be accessible to external entities. Moving into cloud environments can provide opportunities to change this approach to security as the accessibility of these environments can make it easy to inadvertently make resources available outside their intended scope. Additionally, organizations commonly make use of the ease of deployment of new resources provided by cloud environments to enable a broader set of agency users to create new services and infrastructure. Without appropriate security controls, these users may deploy resources without understanding the security implications, and potentially introducing vulnerabilities into the agency's deployed environment. Agencies will need to understand the security controls available for CSP environments that can enable automatic detection and/or prevention of deployment configurations that are out of alignment with agency security policy.

As part of migrating to cloud environments, agencies can consider strong configuration management practices like Development, Security, and Operations (DevSecOps) that help minimize opportunities for misconfiguration by automating and integrating security into the deployment process, and potentially enabling the use of cloud-native solutions to maintain compliance and monitor for vulnerabilities or malicious activity. While these can reduce the potential for misconfiguration, problems in the developed configurations, flaws in the development or deployment pipeline, and issues in the software supply chain can still introduce misconfigurations and vulnerabilities into the deployed infrastructure, and agencies should ensure that their vulnerability management procedures account for potential misconfiguration. Tools that automate the detection and remediation of these types of vulnerabilities can help ensure ongoing security.

4.2.2.7 Changes to Deployment, Maintenance, and Operation

Agencies are deploying services in the cloud as part of their modernization efforts. To effectively make use of cloud environments, agencies may need to look beyond simply lifting and shifting applications, but fundamentally changing how applications are developed, deployed, secured, and operated. There are a wide range of CSPs, cloud native tools, and ways to deploy applications, and agencies will need to understand the variety of service options and how their policies and procedures should be updated to account for their service selections. The changes to policies or procedures will affect the entire workflow as applications are

transitioned, from planning for migration, to developing and testing applications, to planning for maintenance, operations, and ongoing cybersecurity. Agency cloud deployments (SaaS, PaaS, and IaaS) may also integrate services from multiple CSPs and multiple regions, which can be a very different model from an agency providing all services on-premises. Agencies will need to perform due diligence while updating their policies and procedures for development, deployment, maintenance, and operation to ensure their deployments meet agency cybersecurity risk and security objectives.

4.2.2.8 Cloud Security Management Solutions

With the widespread availability of cloud-hosted services, agencies are increasingly using services from multiple cloud vendors, which can introduce additional complexities when managing security. Vendors often include native methods for configuring and securing their cloud environments, but these capabilities may have a limited ability to configure or secure other cloud environments. Agencies' existing enterprise management platforms may have integrated support for their cloud deployments. For agencies looking to deploy new management platforms to centralize management across their cloud environments, there are a variety of services available, commonly falling into a few different categories. Agencies should understand the features and abilities of solutions, as well as their alignment with the agency use case and objectives. Additionally, agencies may need to integrate the solution with their overall development and deployment workflows to ensure alignment of security protections for cloud applications and environments. Many of these cloud security management solutions can facilitate several TIC security capabilities.

- **Cloud Workload Protection Platforms (CWPP):** CWPP can help facilitate visibility and management of security controls in cloud and multi-cloud environments, commonly including functions like system hardening, vulnerability management, host-based segmentation, system integrity monitoring, and application allow lists.
- **Cloud Security Posture Management (CSPM):** CSPM capabilities facilitate monitoring in cloud and multi-cloud environments by identifying, alerting on, and mitigating cloud vulnerabilities. Some CSPM capabilities that focus on managing and securing SaaS applications may be referred to as SaaS Security Posture Management (SSPM) solutions.
- **Cloud Infrastructure Entitlement Management (CIEM):** CIEM capabilities facilitate the management of identities and entitlements in cloud and multi-cloud environments.
- **Cloud-Native Application Protection Platform (CNAPP):** CNAPP capabilities help align the visibility and security protections for deployed cloud applications.

4.2.3 Cloud Connectivity

When using cloud environments, agencies will need to understand the options for connecting their campuses to the cloud environments and the options for enabling access to the deployed resources.

4.2.3.1 Campus Connectivity

Agency campuses have a variety of methods for connecting to cloud environments. Each method presents tradeoffs in terms of flexibility, security, and cost. Independent of the connectivity method chosen, agencies should employ secure transport mechanisms to maintain the confidentiality and integrity of the traffic as it traverses potentially unknown, untrusted, compromised, or shared network infrastructure.

- **Internet:** Agencies can use their existing Wide Area Network (WAN) connectivity to access their cloud environments. This method allows agencies to quickly add new cloud environments without the need for specialized configuration. However, the use of shared WAN connections can have an impact on the latency, performance, or connectivity between the agency and cloud service provider.
- **Private Connection:** Many cloud providers allow for direct connectivity between agency campuses and the cloud provider. These connections need to be specifically configured for each cloud provider, but they can offer improved availability and performance to the cloud environment. While ostensibly

private, agencies should still employ secure transport mechanisms, as the traffic to the cloud provider can still traverse unknown, untrusted, or compromised network infrastructure.

- **Shared Cloud Connection Points:** As an alternative to having individual private connections to each cloud environment, agencies may consider using shared cloud connection points that aggregate direct connections with a variety of cloud providers. These can similarly offer improved availability and performance while easing the effort needed for an agency to connect to a new cloud environment.

4.2.3.2 Service Connectivity

Cloud services can be made available to users and entities through a variety of methods. Different methods offer agencies various levels of control over the devices that access agency services and the layers of protection that can be deployed. Agencies may enable multiple methods for accessing cloud services and environments and may apply different policies depending on the method used.

- **Direct Connection:** Agencies may consider making services available to users directly over the internet, potentially mediated through one or more Security-as-a-Service (SECaaS) protections (e.g., Web Application Firewalls, Next Generation Firewalls, API Gateways). While this deployment model offers the most opportunities for external entities to access the service, it can help ensure uniform security protections by applying the same set of protections independent of where the service is being accessed from.
- **Virtual Private Network:** Agencies may employ VPN or similar secure remote access infrastructure, potentially deployed either in the cloud environment or external to the cloud, and then require users to access cloud services through the VPN. This architecture can provide an additional layer of protection by limiting the accessibility of cloud services. In this architecture, user devices are connected to the cloud environment via the VPN, enabling applications on those devices to potentially access any available resources. To account for this, agencies should consider protections that ensure device compliance while accessing the VPN and protections (e.g., network segmentation, bastion hosts) to limit device access while connected to the VPN. Alternatively, agencies may consider models that more tightly align the connectivity with application workflows (e.g., Secure Access Service Edge (SASE), Zero Trust Network Access (ZTNA)).
- **Remote Desktop Access (RDA):** Agencies may consider using agency-managed desktop instances, potentially deployed either in the cloud environment or external to the cloud. These desktop instances should be made available using secure transport mechanisms (e.g., TLS, VPN, etc.). Agencies may consider protections like gateways or bastion hosts that prevent direct remote access to desktop instances. The use of desktop instance that are deployed and managed by the agency can ensure access and telemetry from a properly configured and managed instance, potentially enabling the agency flexibility in the types of devices users can employ. Agencies can consider treating these desktop instances similarly to remote user devices, aligning the security protections in the cloud deployment, and limiting the trust placed in desktop instances.

4.2.3.3 Management Connectivity

Cloud providers provide interfaces enabling agencies to configure and manage deployments in the provider environments. Often, these consist of web-based user interfaces as well as API interfaces that enable programmatic configuration and management. These management interfaces will form a key component in the operation and maintenance of the agency's cloud deployment, as well as their incident response activities. Given the access that these interfaces can provide, and their importance in enabling the response to potential attacks or compromises of the agency's cloud deployed resources, agencies need to understand the security and availability of these interfaces, and ensure they align with their risk tolerance.

Agencies may use similar methods for accessing these management interfaces with respect to how they make the services themselves available (e.g., Direct Connection, VPN or RDA). However, agencies may also consider different requirements for accessing the management interfaces than for accessing the services themselves. Whichever methods are selected, the agency will need to account for the increased access that management

interfaces can provide and thus apply security protections that align with such increased access and their own risk tolerance.

For IaaS and other cloud deployments, portions of the deployed environment may be configured or managed using agency-defined methods. For example, if an agency builds and deploys a virtual machine into a cloud environment, the cloud provider's configuration and management methods may not have direct visibility or control over the content of the virtual machine or the processes executing within it. In these scenarios, the agency will be responsible for defining the processes for configuring and managing these deployed resources. Often, these processes are executed within the cloud environment itself. However, if these configuration and management methods are made available to users operating outside the cloud environment, the agency will need to ensure that security and connectivity are in line with agency risk tolerances.

4.3 SECURITY PATTERNS

Five security patterns capture the data flows for the IaaS, PaaS, and SaaS guidance in the Cloud Use Case. Each of these has distinct sources, destinations, and options for policy enforcement. Regardless of the options chosen, agencies must ensure they are protecting data in line with their risk tolerances and applicable federal requirements. This is especially important in instances where security policies are being applied by a third party on an agency's behalf or in locations outside the agency's traditional sphere of control. In cases where additional security capabilities are necessary to manage residual risk, agencies should apply the controls or explore options for compensating capabilities that achieve the desired protections to manage risks. The trust levels in these security patterns may not align with agency understanding of their environment and, as such, agencies may determine and label trust zones according to those that best describe their environment. The security patterns include the following trust zones:

- Agency campus
- CSP
- Remote user
- External partner
- External entity
- Web
- Additional Agency Cloud Service

4.3.1 Security Pattern 1: Agency Campus to Cloud Service Provider

Figure 7 illustrates the security pattern where entities within the agency campus trust zone are accessing cloud resources (e.g., through user interfaces, programmatic APIs, etc.). Two options are available for this connectivity and are outlined in Figure 7. Agencies may apply different constraints on connectivity options to different CSP resources. Additionally, agencies may apply different constraints on connectivity for managing the cloud environment than for accessing the cloud resources (e.g., using an API gateway to mediate access to deployed agency services while permitting agency administrators secure direct access to the CSP management interfaces). CSPs may also impose requirements on connectivity. The agency should protect its data in accordance with its risk tolerances and applicable federal requirements.

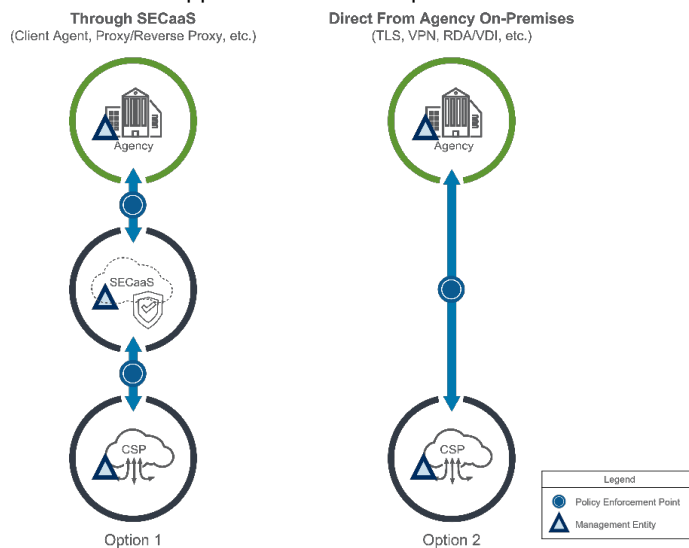



Figure 7: Security Pattern 1 – Agency Campus to Cloud Service Provider



Option 1: The **first option** (left) permits connectivity from on-campus agency entities to cloud resources via a SECaaS provider. Policy enforcement can be performed at the SECaaS provider, the agency campus, and the CSP. Policy enforcement parity between cloud resources can be simplified when all cloud access passes through the same SECaaS provider. Various methods can be used to direct on-campus agency user traffic to the SECaaS provider, including client agents, proxy settings, transparent proxying, and domain name system (DNS). The SECaaS trust zone is labeled with a medium trust level in this option, though agencies may determine and label trust zones according to the trust levels that best describe their environment.

Direct From Agency On-Premises
(TLS, VPN, RDA/VDI, etc.)



Option 2: The **second option** permits connectivity from on-campus agency users directly to cloud resources via protected connections (TLS, VPN, RDA, etc.). Policy enforcement can be performed at the agency campus and the CSP. Policy enforcement parity across multiple campuses can be simplified when policy enforcement is performed at the cloud environment.

4.3.2 Security Pattern 2: Remote User to Cloud Service Provider

Figure 8 illustrates the security pattern where remote agency users are accessing CSP resources (e.g., through user interfaces, programmatic APIs, etc.). Three options are available for this connectivity and are outlined in Figure 8. Agencies may apply different constraints on connectivity options to different CSP resources. Additionally, agencies may apply different constraints on connectivity for managing the cloud environment than for accessing the cloud resources (e.g., using an API gateway to mediate access to deployed agency services while permitting agency administrators secure, direct access to the CSP management interfaces). CSPs may also impose requirements on connectivity. The agency should protect its data in accordance with its risk tolerances and federal requirements.

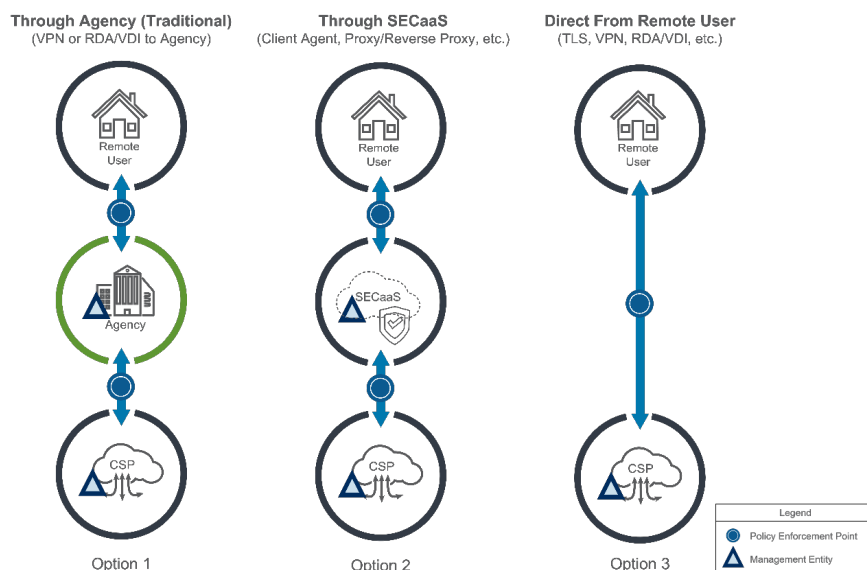


Figure 8: Security Pattern 2 – Remote User to Cloud Service Provider

Through Agency (Traditional)
(VPN or RDA/VDI to Agency)



Option 1: The **first option** (left) aligns with traditional mechanisms for remote users accessing CSP resources. As described in the Remote User Use Case, the remote user establishes a secure connection to the agency campus and accesses the CSP resources through that channel. Policy enforcement can be applied at the agency campus, the CSP and, if possible, on the remote user's device. Policy enforcement parity between remote users and cloud resources can be simplified by applying protections at the agency campus or the CSP.

Through SECaaS
(Client Agent, Proxy/Reverse Proxy, etc.)



Option 2: The **second option** (left) permits connectivity from remote users to cloud resources via a SECaaS provider. Policy enforcement can be performed at the SECaaS provider, the CSP and, if possible, on the remote user's device. Policy enforcement parity between cloud resources can be simplified when all cloud access passes through the same SECaaS provider. Various methods can be used to direct remote user traffic to the SECaaS provider, including client agents, proxy settings, transparent proxying, and DNS. The SECaaS trust zone is labeled with a medium trust level in this option, though agencies may determine and label trust zones according to the trust levels that best describe their environment.

Direct From Remote User
(TLS, VPN, RDA/VDI, etc.)



Option 3

Option 3: The **third option** (left) permits connectivity from remote users directly to cloud resources via protected connections (e.g., TLS, VPN, RDA, etc.). Policy enforcement can be performed at the CSP and, if possible, on the remote user's device. Policy enforcement parity across users can be simplified when policy enforcement is performed at the CSP.

4.3.3 Security Pattern 3: External Entity to Cloud Service Provider

Figure 9 illustrates the security pattern where an external entity (e.g., a public user, an automated external process, an unmanaged IoT device, etc.) can access agency CSP resources (e.g., through user interfaces, programmatic APIs, etc.). With a possibly untrusted entity accessing CSP resources, connections in this security pattern are among the riskiest; therefore, a commensurate amount of rigor should be applied to the security capabilities.

Three options are available for this connectivity and are outlined in Figure 9. Agencies may apply different constraints on connectivity options to different CSP resources. CSPs may also impose requirements on connectivity. The agency should protect its data in accordance with its risk tolerances and applicable federal requirements.

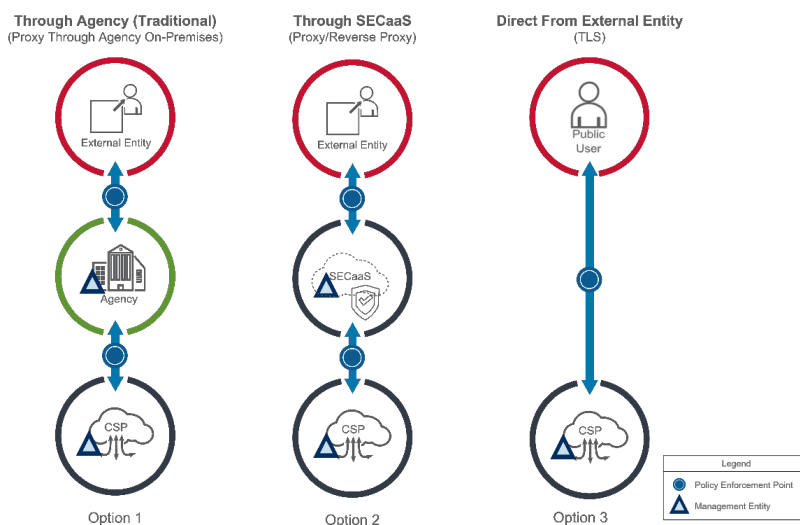


Figure 9: Security Pattern 3 – External Entity to Cloud Service Provider

Through Agency (Traditional)
(Proxy Through Agency On-Premises)



Option 1

Through SECaaS
(Proxy/Reverse Proxy)



Option 2

Option 1: The **first option** (left) aligns with traditional mechanisms for external entities accessing agency CSP resources. The external entity establishes a connection to the agency campus, as described in the Traditional TIC Use Case, and access to the CSP resources is provided through that channel. Policy enforcement can be applied at the agency campus and the CSP. Policy enforcement parity between external entities and cloud resources can be simplified by applying protections at the CSP.

Option 2: The **second option** (left) permits connectivity from external entities to cloud resources via a SECaaS provider. Policy enforcement can be performed at the SECaaS provider and the CSP. Policy enforcement parity between cloud resources can be simplified when all cloud access passes through the same SECaaS provider. Various methods can be used to direct external entity traffic to the SECaaS provider, including DNS and transparent proxying. The SECaaS trust zone is labeled with a medium trust level in this option, though agencies may determine and label trust zones according to the trust levels that best describe their environment.

Direct From External Entity
(TLS)



Option 3: The **third option** (left) option permits connectivity from external entities directly to cloud resources via protected connections (TLS, etc.). Policy enforcement can only be performed at the CSP, which can potentially facilitate policy enforcement parity and resiliency for cloud resources.

4.3.4 Security Pattern 4: External Partners to Cloud Service Provider

Figure 10 illustrates the security pattern where agency CSP resources are provided to an external partner, or agency CSP resources can access resources of an external partner (e.g., through user interfaces, programmatic APIs, etc.). Three options are available for this connectivity and are outlined in Figure 10. Agencies may apply different constraints on connectivity options to different CSP resources or to different external partners. CSPs and external partners may also impose requirements on connectivity. An agency should protect its data in accordance with its risk tolerances and federal requirements.

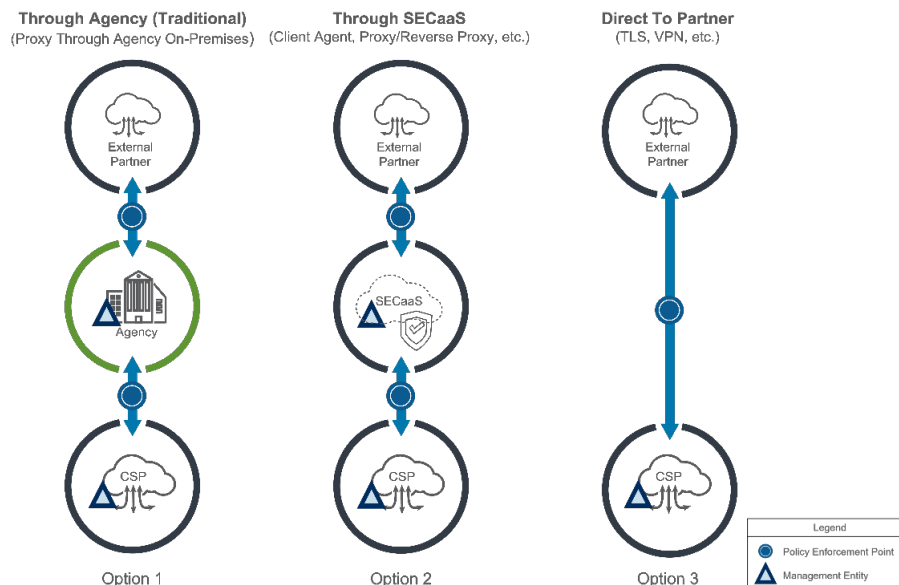


Figure 10: Security Pattern 4 – External Partners to Cloud Service Provider

Through Agency (Traditional)
(Proxy Through Agency On-Premises)



Option 1

Through SECaaS
(Client Agent, Proxy/Reverse Proxy, etc.)



Option 2

Option 1: The **first option** (left) aligns with traditional mechanisms for connectivity between external partners and agency resources, establishing connectivity to the external partner as described in the Traditional TIC Use Case. Policy enforcement can be performed at the agency campus or the CSP. Policy enforcement parity between external partners and agency CSP resources can be simplified by applying protections at the agency campus.

Option 2: The **second option** (left) permits connectivity between external partners and agency CSP resources via a SECaaS provider. Policy enforcement can be performed at the SECaaS provider and the CSP. Policy enforcement parity can be simplified when all connectivity between CSP resources and external partners passes through the same SECaaS provider. Various methods can be used to direct traffic to the SECaaS provider, including proxy settings, DNS, and CSP policy settings. The SECaaS trust zone is labeled with a medium trust level in this option, though agencies may determine and label trust zones according to the trust levels that best describe their environment.



Option 3: The **third option** (left) permits direct connectivity between external partners and agency CSP resources via protected connections (TLS, VPN, etc.). Policy enforcement can only be performed at the CSP, which can potentially facilitate policy enforcement parity and resiliency for cloud resources.

4.3.5 Security Pattern 5: Cloud Service Provider to the Web

Figure 11 illustrates the security pattern where agency CSP resources access resources on the web. Given that agency CSP resources are accessing untrusted resources, connections in this security pattern are among the riskiest; therefore, a commensurate amount of rigor should be applied to the security capabilities.

Three options are available for this connectivity and are outlined in Figure 11. Agencies may apply different constraints on connectivity options to different CSP resources or to different external resources. CSPs may also impose requirements on connectivity. An agency should protect its data in accordance with its risk tolerances and federal requirements.

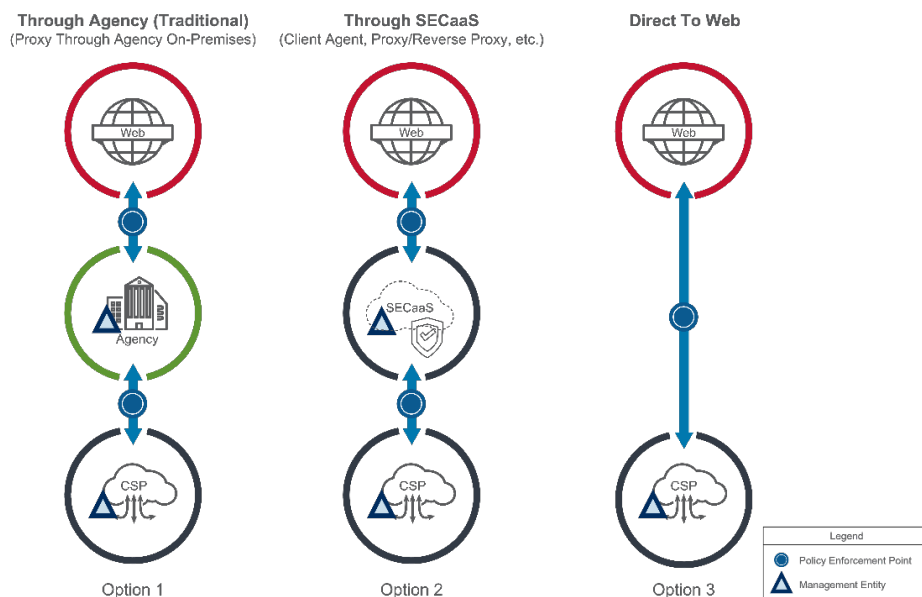


Figure 11: Security Pattern 5 – Cloud Service Provider to the Web

Through Agency (Traditional)
(Proxy Through Agency On-Premises)



Option 1

Through SECaaS
(Client Agent, Proxy/Reverse Proxy, etc.)



Option 2

Option 1: The **first option** (left) aligns with traditional mechanisms for connectivity between agency CSP resources and the internet, with connectivity established as described for an external partner in the Traditional TIC Use Case. Policy enforcement can be performed at the agency campus or the CSP. Policy enforcement parity between agency CSP resources and the web can be simplified by applying protections at the agency campus.

Option 2: The **second option** (left) permits connectivity from agency CSP resources to web resources via a SECaaS provider. Policy enforcement can be performed at the SECaaS provider and the CSP. Policy enforcement parity between agency CSP resources can be simplified when all web access passes through the same SECaaS provider. Various methods can be used to direct agency CSP resource traffic to the SECaaS provider, including client agents, proxy settings, transparent proxying, DNS, and CSP policy features. The SECaaS trust zone is labeled with a medium trust level in this option, though agencies may determine and label trust zones according to the trust levels that best describe their environment.



Option 3: The **third option** (left) permits connectivity from agency CSP resources directly to web resources. Policy enforcement can be performed at the CSP.

4.3.6 Security Pattern 6: Agency Cloud Service to Agency Cloud Service

Figure 12 illustrates the security pattern where agency cloud resources communicate with other agency cloud resources (e.g., through user interfaces, programmatic APIs, etc.). These other agency cloud resources could be hosted in different cloud environments and could have differing deployment models (including multi-region, -tenant and -cloud) and service models (IaaS, SaaS, etc.), and the communication between these agency cloud resources may traverse a variety of networks, regions or CSPs. These communications may be direct communications between agency deployed resources, but could include management traffic where a service or entity in one environment is responsible for managing a service in another environment or the other environment itself. Three options are available for this connectivity and are outlined in Figure 12. Agencies may apply different constraints on connectivity options to different CSP resources or between different CSP resources. Additionally, agencies may apply different constraints on connectivity for managing the cloud environment than for accessing the cloud resources (e.g., using an API gateway to mediate access to deployed agency services while permitting agency administrators secure, direct access to the CSP management interfaces). CSPs may also impose requirements on connectivity. An agency should protect its data in accordance with its risk tolerances and federal requirements.

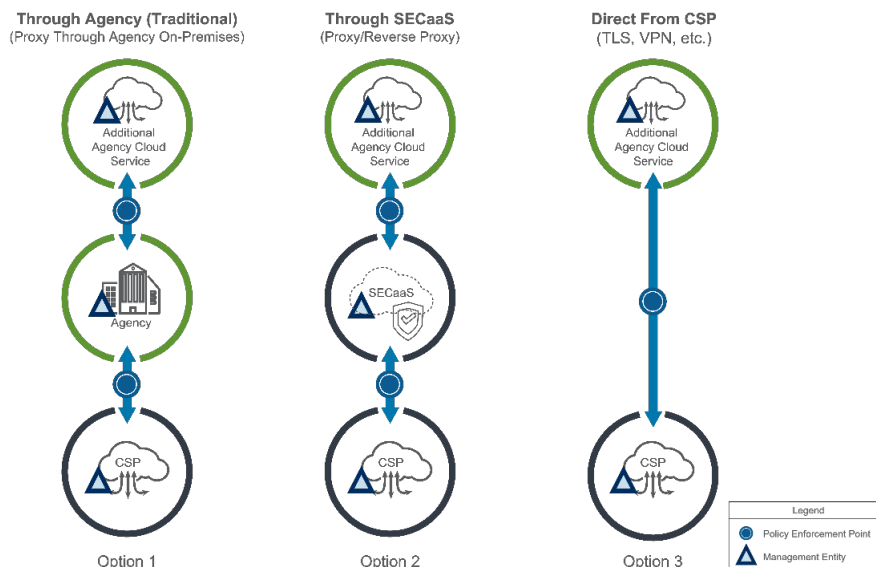


Figure 12: Security Pattern 6 – Cloud Service Provider to Additional Agency Cloud Service

Through Agency (Traditional)
(Proxy Through Agency On-Premises)



Option 1: The first option (left) aligns with traditional mechanisms for connectivity between distributed agency cloud resources, establishing connectivity through the agency campus as described in the Traditional TIC Use Case. Policy enforcement can be performed at the agency campus or the cloud resources. Policy enforcement parity between cloud resources can be simplified by applying protections at the agency campus.

Through SECaaS
(Proxy/Reverse Proxy)



Option 2

Option 2: The second option (left) permits connectivity between Agency cloud resources via SECaaS. Policy enforcement can be performed at the SECaaS or the cloud resources. Policy enforcement parity between cloud resources can be simplified when all access passes through the same SECaaS. Various methods can be used to direct traffic to the SECaaS, including cloud configuration, proxy settings, transparent proxying, and DNS. The SECaaS trust zone is labeled with a medium trust level in this option, though agencies may determine and label trust zones according to the trust levels that best describe their environment.

Direct From CSP
(TLS, VPN, etc.)



Option 3

Option 3: The third option (left) permits direct connectivity between agency cloud resources via protected connections (TLS, VPN, etc.). Policy enforcement can be performed at either cloud resource.

4.4 APPLICABLE SECURITY CAPABILITIES

The Security Capabilities Catalog¹³ contains a table of universal and PEP security capabilities that apply across use cases, but not all apply to every use case. Each use case will contain a set of relevant security capabilities, based on agency pilot implementations and best practices. Additional security capabilities may be employed by agencies to reflect agency requirements, risk tolerances, and other factors. The IaaS, PaaS, and SaaS guidance in the Cloud Use Case is one use case where some PEP security capabilities are not applicable.

¹³ Trusted Internet Connections webpage TIC Guidance Quick Links under Security Capabilities Catalog, Cybersecurity and Infrastructure Security Agency. <https://www.cisa.gov/tic>.

For traceability, the security capabilities not included in this section of the use case are listed below by PEP capability group. Of note, the Email capabilities are not included here as they are part of the Email-as-a-Service Use Case covered in Section 5.

- Email: All
- Networking: Host Containment
- Unified Communication and Collaboration: All

Due to the unique security considerations for the Cloud Use Case, new security capabilities are included. Of note, two new PEP groups have been added: Services and Identity. These capabilities may be added to the next version of the Security Capabilities Catalog upon finalization of this use case. The new security capabilities are detailed in the subsequent tables and listed here by PEP capability group for traceability.

- Universal: Supply Chain Risk Management
- Universal: Resource Lifecycle Management
- Universal: Security Test and Exercise
- Universal: Continuous Monitoring Reporting
- Universal: Governance and Policy Auditing
- Networking: Resource Containment
- Domain Name System: CISA's Protective DNS Service
- Intrusion Detection: Network Detection and Response
- Enterprise: Costs Monitoring
- Data: Data Labeling
- Data: Data Inventory
- Services: All
- Identity: All

Finally, because this section of the use case combines SaaS, PaaS, and IaaS, the universal and PEP security capability guidance in this section includes general guidance applicable to any cloud deployment, as well as specific guidance that is unique to one or more deployment models, as needed.

The universal and PEP security capability guidance in this section includes general guidance applicable to any cloud deployment, as well as specific guidance that is unique to one or more deployment models, as needed.

4.4.1 Universal Security Capabilities

Universal security capabilities are enterprise-level capabilities that outline guiding principles for TIC use cases and apply across all use cases. Agencies have the discretion to determine the level of rigor necessary for applying universal security capabilities in accordance with federal guidelines and their risk tolerance. Universal security capabilities will be considered differently across SaaS, PaaS, and IaaS cloud deployments.

In general, given the shared security model presented in Section 4.2.1, agencies will have less control in SaaS than in PaaS or IaaS. Thus, for most of the universal security capabilities, agencies must understand what is provided by vendors, what is required of the agency, and how to integrate capabilities across multiple CSPs to have an enterprise solution to fulfilling each capability.

Table 2 provides a list of the universal security capabilities that apply to the IaaS, PaaS, and SaaS guidance in the Cloud Use Case and implementation guidance for agencies to consider. Most agencies will have an existing enterprise solution for the universal security capabilities; as agencies deploy the IaaS, PaaS, and SaaS guidance in the Cloud Use Case, the guidance below can be integrated into their existing solutions. While

universal security capabilities are broadly applicable, the circumstances and threats associated with cloud require agencies to consider the security challenges that may need to be addressed.

Table 2: Universal Security Capabilities for IaaS, PaaS, and SaaS

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
Backup and Recovery	Backup and recovery entails keeping copies of configuration and data, as needed, to allow for the quick restoration of service in the event of malicious incidents, system failures, or corruption.	<p>Agencies should ensure that cloud configuration and data are backed up in accordance with agency risk tolerance and applicable federal requirements. Agencies should consider storing backups in separate geographic regions to enable restoration if a region becomes unavailable. Agencies may consider storing cloud backups in separate cloud tenancies or in other cloud environments, using least privilege and separation of duties to limit the potential exposure of the backups. Agencies may also consider keeping copies of the backups in locations outside the cloud environment to ensure their accessibility in the event of the unavailability or compromise of the agency cloud environment.</p> <p>Backups should only be stored in secure locations, and should be encrypted while in-transit to the location and while stored at the location. Additionally, agencies need to consider the storage and access of the keys used to decrypt backups to allow for quick recovery while ensuring that a compromise of the backup location cannot compromise the contents of the backups. Backup solutions should be designed (e.g., disconnected differential backups) to allow recovery both from normal failures and attacks such as ransomware.</p> <ul style="list-style-type: none"> • SaaS: Agencies should consider the availability of backup options when selecting SaaS providers, including opportunities to restore service through alternate SaaS providers. • PaaS/IaaS: When feasible, agencies should use technologies that can automate the construction and deployment of cloud environments and application workflows, minimizing the infrastructure that would need to be backed up to be able to restore service. For infrastructure (e.g., virtual machines, specialized containers) that cannot be easily reconstituted, agencies should perform regular backups.
Central Log Management with Analysis	Central log management with analysis is the collection, storage, and analysis of telemetry, where the collection and storage are designed to facilitate data fusion and where the security analysis aids in discovery and response to malicious activity.	<p>Agencies need to understand their visibility in the cloud environment, including the effect of service offerings on level of visibility or retention time for telemetry. Agencies may, where possible, tailor the retention times to account for risk tolerance, auditing requirements, storage capacity, incident response requirements, and agency need. Agencies should consider how best to integrate cloud telemetry into their overall log management and security analysis solutions, potentially ingesting the cloud telemetry into a centralized collection and storage location. The integrated approach combines security-relevant logging information collected from multiple components within the CSP, and possibly across CSPs. Aggregation of myriad data streams is generally accomplished with a centralized log aggregator and filtering system. The agency can apply artificial intelligence (AI) and machine learning (ML) techniques for heuristic-based anomaly detection, threat and advanced persistent threat detection, and risk and compliance assessment analysis. For telemetry data kept in the cloud environment, agencies need to ensure the available retention times meet their needs and federal requirements,¹⁴ and that their security analysis and incident response workflows can integrate and account for any differences in telemetry availability or retention. For all cloud telemetry, agencies need to account for the possibility for subversion of the collection or availability of telemetry.</p>

¹⁴ "Improving the Federal Governments Investigative and Remediation Capabilities Related to Cybersecurity Incidents," Office of Management and Budget M-21-31 (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>.

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
Configuration Management	Configuration management is the implementation of a formal plan for documenting and managing changes to the environment and monitoring for deviations, preferably automated.	Agencies should consider how best to integrate cloud deployments into their overall configuration management solution, including potential opportunities for orchestration, change control, and reversion to a known good state. As the choice of cloud provider or cloud deployment model can impact the opportunities for integration, agencies may consider the opportunities for integration early in the procurement process. Agencies should consider the use of development and deployment practices, like DevSecOps, that automate and orchestrate the deployment, maintenance, and security of their cloud environments. Agencies may consider the use of Infrastructure-as-Code (IaC) deployment models, potentially using cloud-native solutions, to enable the integration of the cloud environment into their development processes. Agencies may consider tools and capabilities, like CWPP and CNAPP, that can integrate an understanding of the deployed applications and environments and tailor their security protections accordingly. Agencies need to understand how the security of their development and deployment practices affects their cloud environments to ensure the security of the end-to-end deployment lifecycle.
Incident Response Plan and Incident Handling	Incident response planning and incident handling is the documentation and implementation of a set of instructions, procedures, or technical capabilities to sense and detect, respond to, and limit consequences of malicious cyberattacks, and to restore the integrity of the network and associated systems.	<p>Incident response is shared responsibility of the agency and CSP. In general, there will be a gradient of incident response capabilities provided, depending on the service offering type (SaaS, PaaS, or IaaS). In general, in SaaS, agencies will have less visibility and rely on the CSP for incident handling. Agencies should be aware of what incident response capabilities are provided with a service offering, and how they will be notified in the event of a cybersecurity incident that affects the application, or underlying operating system, networks, and hardware.</p> <p>Agencies should update any incident response plans as environments and applications are deployed in the cloud.¹⁵ Agencies should recognize and understand the differences and challenges associated with incident response and handling in the cloud, including lack of access to physical hardware. Cloud environments may have more ephemeral resources than traditional on-premises deployments, and agencies will need to ensure that sufficient telemetry and logging is available to enable accurate reconstruction of incidents that involve ephemeral resources. Incident response plans should consider when a cloud environment has an outage. Agencies should monitor cloud services for misuse or breach and adapt response plans and activities accordingly. Agencies should consider deploying native CSP and third-party tools for incident response. Agencies should evaluate each CSP for its incident response capabilities and integrate this into its incident response plan and handling. Agency response plan should include how the agency will coordinate and collaborate with CSPs for prompt and effective response.</p> <ul style="list-style-type: none"> • SaaS: Agencies should be aware of how they will be notified in the event of a cybersecurity incident that affects the application or underlying operating system, networks, and hardware. • PaaS/IaaS: In a PaaS or IaaS environment, agencies have more responsibility for incident handling; however, in the event of an incident, agencies must rely on the CSP for access to physical networks and hardware.
Inventory	Inventory entails developing, documenting, and maintaining a	For on-premises computing, inventory involves documenting physical assets. As agencies move to cloud environments, this creates new considerations and opportunities for managing and tracking agency cloud assets. Cloud assets include compute resources (e.g., virtual machines, servers, containers,

¹⁵ "Cybersecurity Incident and Vulnerability Response Playbooks," Cybersecurity and Infrastructure Security Agency (2021). https://www.cisa.gov/sites/default/files/publications/Federal_Government_Cybersecurity_Incident_and_Vulnerability_Response_Playbooks_508C.pdf.

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
	current inventory of all systems, networks, and components so that only authorized devices are given access, and unauthorized and unmanaged devices are found and restricted from gaining access.	<p>Function-as-a-Service), storage resources (e.g., block storage, file storage, databases), and other platform assets (e.g., message queues, API gateways, Web Application Firewalls). With malicious entities commonly moving laterally among agency environments, whether on-premises or in the cloud, agencies need to have a strong understanding of all the resources they have deployed.</p> <p>Most CSPs provide dashboards or APIs for tracking these assets and for obtaining current, and often historical, information about deployed cloud resources. While these tools can make it straightforward to track resources in a single cloud vendor, it can be difficult to build a holistic view across all agency assets, both on-premises and in the cloud. Integrating resources from multiple cloud providers can be even more difficult because vendors may use different names, have different properties about a given type of resource, or may even have entirely different types of resources. Agencies need to account for these differences and integrate them into an enterprise inventory.</p> <p>Inventory also involves network asset tracking, including connections into and out of agency environments. Cloud environments increase these connection points, and agencies may also need to track all methods available to access agency resources from entities outside the cloud environment (e.g., VPN, RDA, or direct connectivity to agency resources), but may also need to account for the accessibility of agency resources to other tenants in the cloud environment.</p>
Least Privilege	Least privilege is a design principle whereby each entity is granted the minimum system resources and authorizations that the entity needs to perform its function.	<p>Agencies should maintain visibility into the permissions and their use across their cloud environments to enable enterprise-wide application of least privilege principles, and to allow for the identification and removal of over-provisioned or inactive permissions. Agencies need to consider permissions enterprise wide, including on-premises, cloud, application, and data permissions. Agencies should consider methods for ensuring continuous compliance of least privilege across cloud environments. Agencies should ensure that permissions, especially those with a potential for abuse, apply only during the necessary duration and, when feasible, that users or entities employ on-demand methods for enabling those permissions only for specific resources and only for the time necessary to perform activities on those resources. Agencies may also consider separation of duties for the of performance sensitive activities to decrease the potential for abuse. Agencies should account for device security posture and compliance, and anomalous or suspicious login or user behavior when applying least privilege controls. Agencies may consider using CIEM tools to manage least privilege enterprise wide.</p>
Secure Administration	Secure administration entails securely performing administrative tasks, using secure protocols.	<p>Agencies should ensure only secure protocols can be used to perform administrative functions. Where necessary, agencies should disable all forms of access via insecure means but may deploy compensating protections that ensure availability only via secure protocols. Access to administrative functions should only be available after authentication via strong mechanisms, like phishing-resistant MFA, and should, when possible, integrate additional information like users' device posture before granting access. The access to administrative functions should only be provided for the duration necessary to perform the function and, when feasible, on-demand authentication and authorization methods should be required to enable that access. For administrative functions performed by automated processes, agencies should consider the use of cloud-native tools to manage credentials by taking advantage of one-time passwords, regular key schedules, and other similar technologies to minimize the risk of a leaked credential.</p>

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
		<p>Agencies should strongly apply least privilege to administrative functions and should consider employing separation of duties to ensure that no single account has complete administrative access to cloud environments. Agencies should consider the creation of global administrator accounts that are only to be used in emergencies, commonly referred to as “Break Glass” accounts, to gain administrative access to the cloud environment. These emergency accounts should be well-protected, and agencies should consider requiring the coordination of multiple agency users to enable access to the account. Agencies should enable extensive logging and auditing of administrative activities and consider capabilities that detect anomalous administrative activities. Agencies should consider the potential for administrative accounts to disable or prevent access to alerts or logs when determining how to handle logging and alerting.</p>
Strong Authentication	Strong authentication verifies the identity of users, devices, or other entities through rigorous means (e.g., MFA) before granting access.	<p>Agencies should, wherever possible, employ phishing-resistant MFA¹⁶ as part of verifying identity.¹⁷ Agencies should consider tools like behavioral baselining to detect when a user’s behavior deviates from norms and adaptive authentication, which allows policy to require stricter confirmation of identity when more sensitive access is requested. Agencies should also consider the security posture and compliance of devices that agency users use to access cloud resources.</p> <p>Cloud environments often have tokens, like API keys, that can be used to enable access to cloud services potentially without a commensurate level of identity verification. Agencies will need to understand and track the creation and lifecycle of these tokens.</p>
Time Synchronization	Time synchronization is the coordination of system (e.g., servers, workstations, network devices) clocks to minimize the difference between system clock times and enable accurate comparison of timestamps between systems.	<p>Agencies should understand the synchronization for cloud telemetry generated by services in the cloud environment, and they should account for that when integrating cloud telemetry with telemetry from other cloud or on-premises environments.</p> <p>Where time synchronization options can be configured, agencies may consider the use of agency campus time sources or other external authoritative sources. When choosing external time servers, agencies should understand and account for the effect that device stratum tolerances, latency, link reliability, and other factors can have on time synchronization. Additionally, agencies should ensure appropriate security controls are in place to mitigate opportunities for actors to manipulate the time for agency cloud deployments or otherwise interfere with time synchronization.</p>
Vulnerability Management	Vulnerability management is the practice of proactively working to discover vulnerabilities by including the use of both active and passive means of discovery and by taking action to	<p>The addition of cloud environments can represent a substantial change to an agency’s enterprise and agencies will need to ensure that their overall vulnerability policies and procedures account for these cloud environments. This will ensure that vulnerabilities are being managed in a holistic manner across the agency enterprise.</p> <p>Agencies may choose to apply their existing vulnerability management solutions to their cloud environments. However, agencies may also consider new solutions that align more directly with a given cloud environment. Where multiple solutions are being used, the agency needs to understand the differences between them to ensure an accurate understanding of their</p>

¹⁶ CISA Phishing Resistant Multi-factor Authentication, www.cisa.gov/mfa.

¹⁷ “Moving the U.S. Government Toward Zero Trust Cybersecurity Principles,” Office of Management and Budget M-22-09 (2022). <https://www.whitehouse.gov/wp-content/uploads/2022/01/M-22-09.pdf>.

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
	mitigate discovered vulnerabilities.	<p>overall vulnerability management process, including any limitations. Agencies may need to account for any limitations that cloud environments place on the types of vulnerability discovery and mitigation solutions that can be used for the environment or resources in the environment.</p> <p>Cloud environments offer a wider variety of resources than a traditional deployment, and the types of vulnerabilities and the mitigations available in a cloud environment may differ from traditional mitigations. For example, resources that allow for elastic expansion may be vulnerable to attacks that increase the costs incurred instead of a traditional denial-of-service.</p> <p>Understanding and mitigating these vulnerabilities may not be covered by traditional vulnerability management tools. Agencies should understand these differences in the vulnerabilities applicable to differing cloud resources, as well as the mitigations for these vulnerabilities. Additionally, agencies need to understand whether their vulnerability management solutions can detect these vulnerabilities.^{18 19}</p>
Patch Management	Patch management is the identification, acquisition, installation, and verification of patches for products and systems.	<p>Agencies may have limited visibility into the systems that comprise the cloud environment, limiting their ability to identify needed patches or verify their application. Agencies need to understand the guarantees that cloud providers make toward patching. Additionally, agencies need to understand how the patch procedures can affect agency cloud resources to ensure their cloud-deployed resources can be resilient as cloud providers patch their systems.</p> <p>For agency-deployed components, agencies need to consider how best to integrate the components into their overall patch strategy. They may consider applying their existing patching strategies directly to the cloud-deployed resources. Alternatively, agencies may consider integrating patch management more holistically into the development and deployment processes for their cloud deployments to enable quick reconstitution of cloud resources with the appropriate patches applied.</p> <p>When agency patch procedures use on-premises patch repositories, agencies will need to account for the update procedure in situations where a cloud provider loses connectivity to the agency on-premises location. Agencies should account for the resiliency of the resource during the patching process, including ensuring the ability to revert to a known good state in case a patch creates problems for the cloud deployment.²⁰</p>
Auditing and Accounting	Auditing and accounting include capturing business records (e.g., logs and other telemetry), making them available for auditing and accounting as required, and	Agencies should ensure that their auditing of cloud service activity and business records (including billing) aligns with agency requirements and risk tolerance. Agencies should work to integrate the records generated in the cloud environment into their existing auditing and accounting solutions to enable enterprise-wide visibility. Agencies should understand the visibility available in the cloud environment, as well as how tiers of service can affect that visibility, to align the level of visibility with agency requirements and risk tolerance. Agencies should understand their expected resource usage and monitor for anomalous usage. Furthermore, agencies should enforce more detailed audit logging for their high-risk cloud deployments.

¹⁸ "Binding Operational Directive 22-01: Reducing the Significant Risk of Known Exploited Vulnerabilities," Cybersecurity and Infrastructure Security Agency (2021). <https://cyber.dhs.gov/bod/22-01/>.

¹⁹ "Known Exploited Vulnerabilities Catalog," Cybersecurity and Infrastructure Security Agency (2022). <https://www.cisa.gov/known-exploited-vulnerabilities-catalog>.

²⁰ "Binding Operational Directive 22-01: Reducing the Significant Risk of Known Exploited Vulnerabilities," Cybersecurity and Infrastructure Security Agency (2021). <https://cyber.dhs.gov/bod/22-01/>.

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
	designing an auditing system that considers insider threat (e.g., separation of duties violation tracking) such that insider abuse or misuse can be detected.	
Resilience	Resilience entails ensuring that systems, services, and protections maintain acceptable performance under adverse conditions.	Agencies should consider capabilities (e.g., distributed denial-of-service (DDoS) protection, elastic expansion, and content delivery network (CDN)), that help facilitate resilience for their cloud environments. These capabilities may be native to the cloud environment or may be deployed through external providers. When agencies employ cloud environments, they should account for the resiliency of the network connectivity between their agency campuses and the cloud environments. For cloud environments where the agency does not have direct visibility or control over the resiliency services offered by the cloud environment, they need to understand the services protecting the cloud environment, any SLAs governing the provision of those services, and whether the level of resilience aligns with agency need.
Enterprise Threat Intelligence	Enterprise threat intelligence is the usage of threat intelligence from private or government sources to implement mitigations for the identified risks.	Agencies should sufficiently understand the threats to their cloud environments to align their threat intelligence feeds ^{21,22} with those threats. Agencies may need to augment their existing feeds with additional feeds to ensure a commensurate level of protection. Additionally, agencies should understand whether the security capabilities being deployed to protect their cloud environments can integrate intelligence threat feeds. If the capabilities are not able to use existing intelligence threat feeds, the agency should understand which feeds the capabilities is ingesting and the differences between those feeds and the existing agency solution.
Situational Awareness	Situational awareness is maintaining effective current and historical awareness across all components.	<p>Agencies should integrate their cloud environments into their overall situational awareness solutions to ensure enterprise-wide visibility. Agencies should understand how tiers of service can affect their visibility into the cloud environment and the telemetry available from it. Agencies may consider utilizing cloud-native methods that enable visibility into the cloud environment but need to account for the potential for increased complexity in workflows for using and integrating the information into an enterprise-wide visibility. For agencies integrating cloud environment telemetry into their existing systems and workflows, they will need to account for the accessibility of telemetry in a way that enables integration, as well as differences in the types and makeup of the telemetry available from the cloud environment.</p> <p>To ensure an accurate understanding of the cloud environment, especially when visibility is limited, agencies may need to integrate information provided by the cloud provider detailing their activities that the agency does not have direct visibility into, including environment changes, security threats, roadmaps, etc.</p>
Dynamic Threat Discovery	Dynamic threat discovery is the practice of using dynamic	Agencies should consider solutions for baselining, heuristics, and threat detection that can directly integrate and analyze cloud environment telemetry. When agencies have multiple cloud environments, they may consider separate solutions to align most effectively with each cloud

²¹ "Service Models for Cyber Threat Intelligence White Paper," Cybersecurity and Infrastructure Security Agency (2021). <https://www.cisa.gov/publication/service-models-cyber-threat-intelligence-white-paper>.

²² "Automated Indicator Sharing," Cybersecurity and Infrastructure Security Agency. <https://www.cisa.gov/ais>.

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
	approaches (e.g., heuristics, baselining) to discover new malicious activity.	environment. However, agencies should have enterprise-wide visibility across all agency environments. Agencies need to account for how service levels or deployments can affect the available telemetry and the ability to discover malicious activity. Agencies should, when possible, ensure that dynamic threat discovery solutions can integrate user device, location, and network information, as well as application-level logs and data usage telemetry. This will help provide a broad understanding of user and entity behavior.
Policy Enforcement Parity	Policy enforcement parity entails consistently applying security protections and other policies, independent of the communication mechanism, forwarding path, or endpoints used.	When agency services are available via multiple conveyance methods (e.g., private connection, direct from the internet), agencies should consider protections that apply independent of conveyance mechanisms, potentially focusing on protections more closely aligned with the applications or data to provide consistency. Agencies may consider aligning protections according to user roles, device security posture and compliance, and anomalous or suspicious login or user behavior.
Effective Use of Shared Services	Effective use of shared services means that shared services are employed, where applicable, and individually tailored and measured to independently validate service conformance, and offer effective protections for tenants against malicious actors, both external and internal to the service provider.	Agencies can utilize CSP-provided shared services to facilitate administration, development, operation, and security activities. While agencies may be able to “lift-and-shift” their existing applications, workflows and processes into cloud environments, there may be opportunities for more effective use of the shared services available from CSPs through the adoption of cloud-native workflows and processes. Agencies will need to understand the variety of services available for their chosen CSPs, and how their workflows and processes can be updated to best utilize these services. Agencies may need to look more holistically at their overall enterprise to discover opportunities where shared services can be used to facilitate or improve agency workflows and processes.
Integrated Desktop, Mobile, and Remote Policies	Integrated desktop, mobile, and remote policies define and enforce policies that apply to a given agency entity independent of its location.	If an agency is using a Desktop-as-a-Service (DaaS) offering to deliver virtual desktops to agency users, agencies should maintain security parity across policies for DaaS and other devices. This will help provide consistent protection and minimize user workarounds that could bypass desired security. Agencies may be able to apply their existing policy mechanisms to the DaaS instances, enabling common management of user policies. However, the existing policy mechanisms may require controls or capabilities unavailable in the DaaS environment or may not be well-aligned with the DaaS. In these scenarios, agencies need to understand the controls and capabilities that are offered, how they compare to the existing controls and capabilities, and how to align them to ensure commensurate security or security that aligns with the risks and threats associated with the cloud or DaaS environments.
User Awareness and Training	User awareness and training entails that all users are informed of their roles and responsibilities, and	If a user’s roles or responsibilities change because of a cloud deployment, the agency should make users aware of how these changes affect their cybersecurity responsibilities. Users should be trained to interact securely in new environments.

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
	that appropriate cybersecurity education is provisioned to enable users to perform their duties in a secure manner.	The cloud migration team, cloud developers, cloud administrators, security architects, incident response teams, and related IT staff should have the necessary training to support all agency cloud services. Analysts may need to be trained to understand new resources and environments. These professionals may need to be trained to support cloud services at several CSPs. These types of highly specialized training courses are typically provided by CSPs. Agencies should consider refresher training for both users and IT staff as cloud technology advances and new cybersecurity threats are discovered.
Supply Chain Risk Management	Supply chain risk management involves implementing a systematic process for managing cyber supply chain risk exposures, threats, and vulnerabilities throughout the supply chain and developing risk response strategies to the risks presented by the supplier, the supplied products and services, or the supply chain.	<p>When agencies acquire services from CSPs, they should specifically consider and implement supply chain risk management as part of their existing risk management activities. FedRAMP has established cloud services security guidelines as a standardized approach to assessing and authorizing cloud products and services. Also, NIST Special Publication (SP) 800-161²³ provides additional guidance for agencies implementing supply chain risk management.</p> <ul style="list-style-type: none"> • PaaS/laaS: As agencies develop and deploy applications and services in the cloud, they should consider the supply chain of any third-party products or services used. The considerations should be similar to how they assess and consider supply chain for products used on-premises. • SaaS: Often a SaaS product is provided by a single vendor. However, when agencies add a third-party service to a SaaS deployment, they should consider the supply chain of that service.
Resource Lifecycle Management	Resource lifecycle management is the end-to-end process of managing resources from development to operation to retirement, such that resources are provisioned and decommissioned in conjunction with the applications they support.	<p>While traditional environments often track the lifecycle of physical devices and applications, cloud environments can provide a variety of types of resources that agencies may need to track as part of the deployment and decommission process. Agencies need to understand how to integrate the deployment, tracking, and removal of cloud resources into their overall lifecycle management workflows.</p> <p>The accessibility of cloud resources and the potential for reuse of cloud resources can make it imperative for agencies to ensure the removal of resources that are no longer in use. Agencies should consider solutions that integrate the deployment process into the overall development process to ensure cloud resources align with the deployed applications.</p>
Security Test and Exercise	Security tests (e.g., penetration testing or red teaming) verify the extent to which a system resists active attempts to compromise its security. Security exercises are	Agencies need to understand the policies that each cloud provider has concerning security testing to align agency testing and exercise procedures, including ensuring that the security testing is in accordance with CSP acceptable testing policies. Agencies' security testing and exercising needs to be handled in a holistic pattern, ensuring that each cloud environment is not tested in isolation but as part of an overall defense strategy that accounts for vulnerabilities and attack techniques that may employ multiple environments. Agencies should, where feasible, augment their security testing of cloud environments to include automated security tests that can facilitate appropriate testing as changes occur in the cloud deployment, potentially

²³ "SP 800-161 Cyber Supply Chain Risk Management Practice for Systems and Organizations," National Institute of Standards and Technology (2021). <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-161r1-draft.pdf>.

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
	simulations of emergencies that validate and identify gaps in plans and procedures.	aligning with the security testing in OMB M-22-09. ²⁴ Stronger integration of the Agencies can consider the use of cloud native capabilities to conduct security testing.
Continuous Monitoring Reporting	Continuous monitoring reporting entails the maintenance of ongoing awareness of informational security, vulnerabilities, and threats to support organizational risk management decisions.	Agencies should ensure that they collect and integrate relevant information from the cloud provider as well as their deployment to ensure an accurate understanding of potential vulnerabilities and threats to the deployed services. Agencies should understand what level of visibility they have into that information, and either augment or account for any missing visibility when applying that data. Additionally, the agency will also need to understand the timeliness of that information as well as how long it is retained and ensure that it aligns with agency need.
Governance and Policy Auditing	Governance and policy auditing entails validating the proper definition, application and enforcement of agency rules and policies.	<p>As agencies transition into cloud environments, they will need to understand how to ensure these new deployments are compliant with both internal and external governance requirements. Many of these policies are going to need to be deployed and enforced in these new environments, and agencies will need to understand the opportunities for defining policies provided by the cloud provider. Policies should be implemented, where feasible, in ways that enable their automated enforcement. Where policies cannot be deployed or automated through native solutions, there may be opportunities to integrate other capabilities, either directly in the cloud environment or through externally provided solutions, that can be used to implement the policies.</p> <p>As changes occur in the cloud environment, including changes defined by the agency as well as evolution in the solutions offered by the CSP, agencies will need to ensure the continued accuracy of these defined policies and validate their enforcement. Agencies should consider mechanisms that can integrate the policy definition and deployment as part of the development and deployment processes that the agency uses to manage the cloud environment. Agencies should consider validating proper enforcement of these policies as part of their Vulnerability Management strategy.</p>

4.4.2 Policy Enforcement Point Security Capabilities

PEP security capabilities are primarily focused on the network level and inform technical implementation for a given use case, such as communication with agency-sanctioned CSPs. Agencies can implement these capabilities using a variety of methods, including CSP native, third-party provided or agency-deployed solutions. Agencies have the discretion to determine the applicability and level of rigor necessary for applying PEP security capabilities based on the specific cloud service deployed, available policy enforcement options, federal guidelines, and risk tolerance.

From the Security Capabilities Catalog, the PEP security capability groups applicable to the IaaS, PaaS, and SaaS guidance in the Cloud Use Case correspond to the following security functions:

- Files
- Web
- Networking
- Resiliency

²⁴ "Moving the U.S. Government Toward Zero Trust Cybersecurity Principles," Office of Management and Budget M-22-09 (2022). <https://www.whitehouse.gov/wp-content/uploads/2022/01/M-22-09.pdf>.

- DNS
- Intrusion Detection
- Enterprise
- Data Protection
- Identity
- Services

Agencies may determine the applicability and rigor of the security capabilities based on federal guidelines, mission needs, available policy enforcement options, and risk tolerance.

Of note, two new PEP security capability groups have been added: Services and Identity. Security capabilities that are not applicable to this use case are listed at the beginning of Section 4.4. The PEP security capability listing is not exhaustive. Additional security capabilities may be deployed by agencies to reflect their risk tolerances, early adoption of security capabilities, the maturity level of existing cyber programs, etc.

4.4.2.1 Files PEP Security Capabilities

Agencies should employ file protection capabilities to prevent malicious files from being brought into the environments. These protections can help secure deployed cloud resources and help prevent the agency cloud resources from being used to provide malicious files to agency users or other external entities. These capabilities may take the form of solutions to detect malicious files during transmission into or out of the environment. These capabilities could be integrated with the cloud environment or be deployed as part of a SECaaS, CASB or similar solution. Agencies that only apply capabilities to files that are brought into or sent out of the cloud environment need to understand all potential ways that files may be brought into the cloud environment. This is key to ensure enforcement parity across all methods of file ingestion.

When agencies deploy applications into cloud environments, they will need to consider how to ensure that file capabilities are applied to any files sent or received by the deployed applications, and they may need to integrate methods into the applications to enable the quarantine or removal of malicious files. Table 3 lists the applicable Files PEP Security Capabilities for the IaaS, PaaS, and SaaS guidance in the Cloud Use Case.

Table 3: Files PEP Security Capabilities for IaaS, PaaS, and SaaS

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
Anti-malware	Anti-malware protections detect the presence of malicious code and facilitate its quarantine or removal.	<p>Agencies should align anti-malware protections with the potential threats to their cloud environment.</p> <p>For deployments where agencies manage execution environments in the cloud environment (e.g., deployed containers or virtual machines), agencies should ensure that appropriate anti-malware capabilities are applied to deployed execution environments. This will help enable the detection of malware brought into, or executed in, these environments.</p>
Content Disarm and Reconstruction	Content disarm and reconstruction technology detects the presence of unapproved active content and facilitates its removal.	<p>Agencies may consider the use of content disarm and reconstruction technologies to deployments that allow for file submission into the environment, or where entities in the cloud environment may access files from external locations. Content disarm and reconstruction technologies may change documents in ways that render them unsuitable for agency use. In these instances, agencies should consider options for making the original file available to agency users on an as-needed basis. Agencies may also employ methods for agency users to access unmodified files from trusted sources.</p>

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
Detonation Chamber	Detonation chambers facilitate the detection of malicious code using protected and isolated execution environments to analyze the files.	<p>Agencies may consider a centralized detonation chamber capability or a capability deployed to a given cloud environment. When using the centralized model, agencies need to consider how losing connection between the cloud environment and the detonation chamber may affect the security or operation of the cloud deployment. When using a multi-environment model with different detonation chamber capabilities, agencies should understand the differences between the deployed capabilities.</p> <p>Detonation chamber detection capabilities may occur after a file has been ingested into the environment. In these instances, agencies need to understand how files are quarantined or removed from the cloud environments. Agencies need to account for that possibility when building applications.</p>
Data Loss Prevention	Data loss prevention (DLP) technologies detect instances of the exfiltration, either malicious or accidental, of agency data.	<p>Agencies should, where possible, consider data loss prevention in a holistic manner throughout the cloud environment and across the agency enterprise, ensuring data loss can be detected and potentially prevented when occurring over multiple modes of conveyance. Agency data should be brought into or out of cloud deployments using authentication and transport mechanisms that can protect the confidentiality and integrity of the data in line with agency risk tolerances. DLP solutions at agency cloud deployments should ensure that sensitive data does not enter cloud deployments unless it is authorized to be stored in the cloud deployment.</p> <p>DLP solutions at agency cloud deployments must also consider any possibilities for file exfiltration from the cloud. CSPs may offer DLP services that can be added to a service offering. Alternatively, agencies can implement DLP controls.</p> <ul style="list-style-type: none"> • SaaS: SaaS may directly integrate DLP into the service or may support the integration of third-party DLP controls. Agencies must understand what is provided by these offerings and if these solutions are robust enough to detect sensitive agency data. • PaaS/IaaS: DLP may be implemented as part of security controls applying to traffic exiting the cloud deployment, or as controls integrated throughout the deployment, tracking data as it is collected in or moves through the cloud deployment.

4.4.2.2 Web PEP Security Capabilities

The Web PEP capabilities apply when entities in an agency cloud environment accesses the internet, external partners, or other tenants in the same cloud environment. Agencies should, if possible, apply web capabilities to all data flows from entities in the cloud environment to these external resources. These capabilities should be commensurate to those available from the agency campus. Agencies may need to rely on other compensating protections, as cloud environments may not provide the same policy enforcement locations.

Cloud environments can offer numerous avenues for entities to access external environments. Agencies need to understand the various ways entities in the cloud environment can access the internet or other external environments. Agencies also need to understand their visibility and control over this connectivity. Where the agency has limited visibility or control, they should understand what protections the cloud vendor may be applying to those connections and whether they need to apply compensating controls to align that connectivity with their risk tolerances.

Agencies may, where possible, apply the same web protection solution used in their traditional environments to the cloud environment. However, this the solution may need to be augmented for each agency's specific use case, cloud environments, or threats. Alternatively, agencies may consider new solutions that better align with their specific use cases, cloud environments, or threats. When employing a different solution, agencies need to understand the differences between their existing solution and the solution for their cloud environment to maintain a holistic, enterprise wide understanding of their security protections. Table 4 lists the applicable Web PEP Security Capabilities for the IaaS, PaaS, and SaaS guidance in the Cloud Use Case.

Table 4: Web PEP Security Capabilities for IaaS, PaaS, and SaaS

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
Break and Inspect	Break and inspect systems, or encryption proxies, terminate encrypted traffic, logging or performing policy enforcement against the plaintext, and re-encrypting the traffic, if applicable, before transmitting to the final destination.	<p>Agencies may not have the visibility or control necessary to apply traditional break and inspect techniques to many of the traffic flows in cloud environments and may need to investigate alternative security capabilities or deployment architectures to provide commensurate security.</p> <p>Agencies should consider the protections and lifetimes of certificates used as part of Break-and-Inspect certificates. This will decrease the chance of compromise and mitigate the effects of certificate compromise to cloud entities. Agencies may consider the use of cloud-native management tools to manage the certificates.</p>
Active Content Mitigation	Active content mitigation protections detect the presence of unapproved active content and facilitate its removal.	No specific guidance.
Certificate Denylisting	Certificate denylisting protections prevent communication with entities that use a set of known bad certificates.	No specific guidance.
Content Filtering	Content filtering protections detect the presence of unapproved content and facilitate its removal or denial of access.	No specific guidance.
Authenticated Proxy	Authenticated proxies require entities to authenticate with the proxy before making use of it, enabling user, group, and location-aware security controls.	No specific guidance.

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
Data Loss Prevention	DLP technologies detect instances of the exfiltration, either malicious or accidental, of agency data.	<p>Agencies should, where possible, consider data loss prevention in a holistic manner throughout the cloud environment and across the agency enterprise, ensuring data loss can be detected and potentially prevented when occurring over multiple modes of conveyance. Agency data should be brought into or out of cloud deployments using authentication and transport mechanisms that can protect the confidentiality and integrity of the data in line with agency risk tolerances. DLP solutions at agency cloud deployments should ensure that sensitive data does not enter cloud deployments unless it is authorized to be stored in the cloud deployment.</p> <p>DLP solutions at agency cloud deployments must also consider any possibilities for file exfiltration from the cloud. CSPs may offer DLP services that can be added to a service offering. Alternatively, agencies can implement DLP controls.</p> <ul style="list-style-type: none"> • SaaS: SaaS may directly integrate DLP into the service, or may support the integration of third-party DLP controls. Agencies must understand what is provided by these offerings and if these solutions are robust enough to detect sensitive agency data. • PaaS/IaaS: DLP may be implemented as part of egress controls or integrated in controls that apply throughout the deployment.
Domain Resolution Filtering	Domain resolution filtering prevents entities from using unauthorized DNS resolution services over the DNS-over-Hypertext Transfer Protocol Secure (HTTPS), domain resolution protocol.	Depending on the cloud provider or the deployment model, agencies may have limited control over the DNS infrastructure used by services in cloud environments, which may employ the cloud-provided DNS infrastructure. As such, the filtering provided by these protections may be limited to agency cloud resources where the agency has control over the DNS configuration of the resource.
Protocol Compliance Enforcement	Protocol compliance enforcement technologies ensure that traffic complies with protocol definitions, like those documented by the Internet Engineering Task Force (IETF). ²⁵	No specific guidance.

²⁵ "RFCs," Internet Engineering Task Force (2021). <https://www.ietf.org/standards/rfcs/>.

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
Domain Category Filtering	Domain category filtering technologies allow for classes of domains (e.g., banking, medical) to receive a different set of security protections.	No specific guidance.
Domain Reputation Filtering	Domain reputation filtering protections are a form of domain denylisting based on a domain's reputation, as defined by either the agency or an external entity.	No specific guidance.
Bandwidth Control	Bandwidth control technologies allow for limiting the amount of bandwidth used by different classes of domains.	No specific guidance.
Malicious Content Filtering	Malicious content filtering protections detect the presence of malicious content and facilitate its removal.	No specific guidance. ²⁶
Access Control	Access control technologies allow an agency to define policies limiting what actions may be performed by connected users and entities.	No specific guidance.

4.4.2.3 Networking PEP Security Capabilities

Cloud environments and deployments may offer agencies varying levels of visibility and control into traditional networking communication channels. Additionally, cloud environments may have alternative communication channels, including control channels or specialized resource communication mechanisms, that may differ in visibility and control from the environment's traditional networking offerings. Agencies should understand the various methods available for communication in the cloud environment to align security of the environment with their risk tolerance.

The controls for managing communication channels in cloud environments can also differ from traditional environments. To ensure a commensurate set of protections are being applied to agency resources in the cloud, agencies need to understand the controls available and how to best apply these controls to agency use cases. While some environments may enable the use of common traditional networking tools like routers and firewalls, others may only offer higher-level services with little visibility or control at the lower-layer communication channels. Agencies may need to rely on controls for higher-level services (e.g., web application firewalls, CASBs, Software-Defined Perimeters, ICAM access controls) to enable protections similar to those provided in traditional networking environments. Depending on the use case or deployment, agencies may find it easier or more effective to use those higher-level protections even when a cloud environment offers low-level visibility and control. Agencies may also consider methods that combine the network level protections with higher level user or application workload information, allowing a stronger alignment of the network segmentation with microsegmentation principles.

Table 5 lists the applicable Networking PEP Security Capabilities for the IaaS, PaaS, and SaaS guidance in the Cloud Use Case.

²⁶ "Capacity Enhancement Guide on Securing Web Browsers and Defending Against Malvertising for Federal Agencies," Cybersecurity and Infrastructure Security Agency. <https://www.cisa.gov/publication/capacity-enhancement-guides-federal-agencies>.

Table 5: Networking PEP Security Capabilities for IaaS, PaaS, and SaaS

Capability	Description	Use Case and Deployment-Specific Guidance
Access Control	Access control protections prevent the ingest, egress, or transiting of unauthorized network traffic.	Agencies that lack visibility or control over network access protections may consider the use of Software-Defined Perimeter, ICAM ²⁷ or other higher layer access controls. They may also consider using a SECaaS, CASB or similar mechanism to limit access. Agencies may consider the use of cloud-native network access control mechanisms, but may consider, where possible, the deployment of virtualized network access controls.
Internet Address Denylisting	Internet address denylisting protections prevent the ingest or transiting of traffic received from, or destined, to a denylisted internet address.	As many cloud environments apply address denylisting to all traffic incoming to their environment, agencies should align any denylist protections being applied by their cloud environments to their business needs and risk tolerance. If agencies need to supplement the protections provided by the cloud provider, they may consider cloud-native solutions or the use of SECaaS, CASB or similar mechanisms. Agencies may consider, where available, deploying virtualized network denylisting functions to align with agency needs.
Network Segmentation	Network segmentation separates a given network into subnetworks, facilitating security controls between the subnetworks, and decreasing the attack surface of the network.	<p>Agencies may consider, where possible, the use of traditional network segmentation technologies to divide networks. However, agencies may need to use alternative methods for segmenting their network (e.g., isolated environments, Software-Defined Perimeter, or ICAM access controls). Where available, agencies should consider alternative communication channels to ensure proper network segmentation.</p> <p>Cloud environments often provide easier opportunities to integrate microsegmentation. When deploying new infrastructure into cloud environments, agencies should consider opportunities for enabling microsegmentation, instead of high-level network segmentation.</p> <p>When using VPNs or similar technologies to bridge cloud environments with other environments, the agency should ensure, if possible, that the bridged networks are segmented so that least privilege access is maintained, and to limit the scope of a compromise of any environment.</p>
Microsegmentation	Microsegmentation divides the network, either physically or virtually, according to the communication needs of application and data workflows, facilitating security controls to protect the data.	<p>Agencies should consider the use of microsegmentation approaches to network segmentation. In environments where agencies lack visibility or control over the traditional networking mechanisms, agencies may be able to apply microsegmentation techniques through Service Meshes, Software-Defined Perimeter or ICAM access control mechanisms.</p> <p>Agencies should align microsegments with application deployments. Where possible, agencies should consider automated approaches that integrate the microsegment deployment with the application deployment.</p> <p>When using VPNs or similar technologies to bridge cloud environments with other environments, the agency should consider the application microsegmentation techniques to those connections to ensure that only authorized, and potentially authenticated, data flows are permitted between environments.</p>

²⁷ "Federal ICAM Architecture Introduction," General Services Administration (2021). <https://playbooks.idmanagement.gov/arch/>.

Capability	Description	Use Case and Deployment-Specific Guidance
Resource Containment	Resource containment protections enable the removal or quarantine of a resource's access to other resources.	Some resources in cloud environments may be amenable to traditional host containment strategies, especially in situations where the agency has control over the networking and the hosts themselves. However, cloud environments can provide a diverse set of resources with a variety of communication channels and controls to limit access to or from those resources. To prevent compromised resources from accessing other resources, agencies need to understand all of their cloud resources, the communication channels available to those resources, and the opportunities available to restrict access to those resources. For some abstract resources, it may not be feasible to restrict access to any given instance; thus, it may require removal from service for remediation. This is especially true for highly ephemeral abstract resources, like a function-as-a-service, that may only execute in response to a trigger. For resources that are allocated in an automated manner, agencies may consider destroying and recreating the resources, as an alternative to containing the resource. This will allow reversion to a known state, provided that agency understands and can prevent the resource from becoming compromised again.

4.4.2.4 Resiliency PEP Security Capabilities

Resiliency has historically been difficult to achieve in traditional infrastructure. Thus, traditional infrastructures require substantial investments in redundant hardware and locations, which are sometimes seen as underutilized, and create additional configuration complexity. Agencies can simplify resiliency using solutions that enable on-demand scaling and rapid recovery, while regional delivery capabilities can facilitate deployments. Cloud deployments can also be more easily integrated with protections against resource consumption attacks, such as DDoS, whether directly integrated with the cloud deployment or employed through mechanisms like CDNs and CASBs. Agencies can integrate resiliency planning and capabilities into their development and deployment practices to minimize downtime during deployments and migrations, and to facilitate rapid recovery after failures. Table 6 lists the applicable Resiliency PEP Security Capabilities for the IaaS, PaaS, and SaaS guidance in the Cloud Use Case.

Table 6: Resiliency PEP Security Capabilities for IaaS, PaaS, and SaaS

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
Distributed Denial of Service Protections	DDoS protections mitigate the effects of distributed denial of service attacks.	Agencies should employ DDoS protections for their cloud services to ensure the availability of the services, especially for remote users who may need to use traditional networks to access the services. These DDoS protections can be employed either in a cloud-native fashion or potentially through an external provider. For protections provided by external providers, agencies need to ensure that the cloud resource is only accessible through those DDoS protections. Agencies should consider an alternative threat where the DDoS does not render the cloud resource inaccessible but makes use of elastic expansion features to drive up the costs incurred by the agency. Agencies should monitor their cloud resource usage to detect anomalous resource usage.
Elastic Expansion	Elastic expansion enables agencies to dynamically expand	Agencies should understand the opportunities available for elastic expansion in their cloud environments. While some providers may provide

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
	the resources available for services as conditions require.	automated scaling as a feature, others may require the agency to integrate expansion as part of its development and deployment process. In addition to expansion, agencies should consider contraction, or the automatic de-provisioning of resources. If cloud resources are scaled down when demand diminishes, this will reduce an agency's costs and attack surface. As resources are provisioned and/or de-provisioned, these changes should be tracked appropriately, potentially with the agency's inventory system tracking the potential for expansion, tracking expansion, and expansion and resource utilization using monitoring and logging.
Regional Delivery	Regional delivery technologies enable the deployment of agency services across geographically diverse locations.	<p>Agencies should consider regional delivery models to improve the resilience of their services and to reduce the latency for entities accessing the services. Agencies may enable regional delivery through solutions deployed in the cloud environment, or through external services, like CDNs or CASBs.</p> <p>In some cloud environments, agencies may lack the visibility into or control over regional delivery. For these environments, agencies will need to understand the delivery options employed by the provider to ensure they align with agency needs. Where the provided service does not align, agencies may need to augment the service with external services that can enable regional delivery.</p>

4.4.2.5 Domain Name System PEP Security Capabilities

Agencies may not be able to control the DNS resolution used in cloud deployments, and they need to understand and account for any gaps in visibility or protection. Where agencies have control over the DNS resolution used by their cloud deployment, they may consider deploying DNS resolution protections to augment any protections provided by the cloud vendor. Alternatively, agencies may consider having their cloud deployments utilize the agency's existing DNS resolution infrastructure to facilitate centralized DNS security protections while accounting for potential latency or loss of connectivity issues.

Cloud resources can have their own domain names, often hosted in domain hosting services provided by the cloud provider. Agencies need to understand the domain names for their cloud resources and the protections available for those names. Agencies should, where possible, consider integrating the lifecycle of those domain names with their development and deployment processes to facilitate proper functioning and security coverage of the environment. Table 7 lists the applicable DNS PEP Security Capabilities for the IaaS, PaaS, and SaaS guidance in the Cloud Use Case.

Table 7: Domain Name System PEP Security Capabilities for IaaS, PaaS, and SaaS

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
Domain Name Sinkholing	Domain name sinkholing protections are a form of denylisting that protect clients from accessing malicious domains by redirecting the DNS queries for those domains.	Agencies should understand whether the CSP-provided domain resolution service performs domain name sinkholing. If the service does not, agencies should, if possible, consider alternative resolution services that perform sinkholing, or ways to supplement the visibility or protections in the cloud deployment.
Domain Name Verification for Agency Clients	Domain name verification protections ensure that domain name lookups from agency clients, whether for internal or external domains, are validated	Agencies should understand whether the CSP-provided domain resolution service performs DNSSEC. If the service does not, agencies should consider alternative resolution services that perform verification, or ways to supplement the visibility or protections in the cloud deployment.

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
	according to Domain Name System Security Extensions (DNSSEC).	
Domain Name Validation for Agency Domains	Domain name validation protections ensure that all agency domain names are secured using DNSSEC, enabling external entities to validate their resolution to the domain names.	When cloud environments provide domain name hosting, agencies should understand whether those hosted domains are secured using DNSSEC. If not, agencies should ensure that there are DNSSEC-secured domain names available to access the cloud-deployed services and should consider ensuring that those services are only available using the DNSSEC-secured domain names.
Domain Name Monitoring	Domain name monitoring allows agencies to discover the creation of or changes to agency domains.	Agencies should ensure that domain name monitoring solutions integrate domain names created for agency cloud services, including those hosted by the cloud provider. Agencies may consider more tightly integrating an understanding of cloud deployment workflows into the monitoring to detect anomalous domain activity more effectively.
CISA's Protective DNS Service	CISA's Protective DNS Service is a shared service offering that provides domain name sinkholing protections. Agency utilization of this service is mandatory.	Agencies would need to work with CISA ²⁸ to ensure commensurate protections and visibility are available when cloud deployments use domain resolution infrastructure that does not integrate protections from CISA's Protective DNS Service.

4.4.2.6 Intrusion Detection PEP Security Capabilities

Cloud environments can provide agencies with a variety of resource types, which can each have unique risks and potential vulnerabilities. Agencies will need to ensure that intrusion detection solutions account for these differences and provide protections in line with agency needs and risk tolerances. Intrusion detection solutions may need to account for a variety of deployment methods and architectures to handle the diversity of agency cloud needs. As agencies may lack complete visibility and control in cloud environments, intrusion detection will be a shared responsibility. Agencies will need to understand the intrusion detection roles and responsibilities of the vendor to ensure alignment with agency needs and risk tolerances. Agencies also need to understand their visibility into the cloud environments, along with how tier-of-service can affect that visibility, to ensure alignment with their intrusion detection needs. Agencies may deploy intrusion detection solutions in the cloud environment or in external environments. Where solutions are external to the cloud environment, agencies will need to account for how a loss of connectivity between the cloud environment and the intrusion detection solution can affect detection and automated responses. Table 8 lists the applicable Intrusion Detection PEP Security Capabilities for the IaaS, PaaS, and SaaS guidance in the Cloud Use Case.

Table 8: Intrusion Detection PEP Security Capabilities for IaaS, PaaS, and SaaS

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
Endpoint Detection and Response²⁹	Endpoint detection and response (EDR) tools combine endpoint and network event data to	Agencies may consider tailored (potentially native) EDR solutions for cloud deployments, but they need to ensure there is a holistic understanding of the enterprise-wide detection capabilities and endpoint visibility. Agencies should also consider any potential differences in the types of available detections or visibility for individual deployments.

²⁸ OSMO@cisa.dhs.gov

²⁹ Endpoint Detection and Response reference: <https://www.whitehouse.gov/wp-content/uploads/2021/10/M-22-01.pdf>.

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
	aid in the detection of malicious activity.	Agencies will need to consider how to integrate endpoint data from cloud deployments into their overall enterprise-wide situational awareness to enable the commensurate detection of malicious activity, independent of where it occurs in the agency's overall environment. If agencies are considering transitioning from an on-premises to a cloud-based EDR solution, they should account for how the loss of connectivity by the agency campus, branch offices, or remote users might affect detection or response.
Intrusion Detection and Prevention Systems	Intrusion detection systems (IDS) detect and report malicious activity. Intrusion prevention systems (IPS) attempt to stop the activity.	<p>Agencies can leverage cloud native solutions to prevent incidents from occurring in their traditional and cloud deployments. Most CSPs offer advanced AI systems trained on data across their services and customers that outperform traditional AI-detection systems only deployed in one network.</p> <p>Agencies can utilize cloud IPS systems to protect their cloud resources at the network level and prevent attacks from pivoting across cloud resources or from traditional networks into cloud resources.</p>
Adaptive Access Control	Adaptive access control technologies factor in additional context, like security risk, operational needs, and other heuristics when evaluating access control decisions.	Where possible, agencies should determine access to cloud or agency resources according to user roles, device security posture and compliance, and anomalous or suspicious login or user behavior. Additionally, agencies may consider strength of authentication as part of access determination.
Deception Platforms	Deception platform technologies provide decoy environments from individual machines to entire networks, that can be used to deflect attacks on the operational systems supporting agency missions/business functions.	Agencies may consider deception platforms that can be deployed in their cloud environment. These may take the form of honeypot mechanisms that can help detect malicious activities in the cloud environment but may include more advanced decoy network infrastructure that can be used to monitor malicious activity. Agencies need to understand the difference between their existing infrastructure and the new cloud-deployed infrastructure to ensure they align with the threats that the deception platform is targeting. This is especially important for agencies that are integrating cloud-deployed deception environments with their existing deception infrastructure, or agencies that are deploying deception infrastructure in a cloud environment that mimics their other environments.
Certificate Transparency Log Monitoring	Certificate transparency log monitoring allows agencies to discover when new certificates are issued for agency domains.	Agencies should ensure that their certificate transparency log monitoring supports detection for domains used by their cloud deployments. Additionally, agencies need to understand the resources being pointed to by agency domains as IP and internal cloud domain reuse in cloud environments may allow for threat actors to obtain the address or internal cloud domain that an agency domain references.
Network Detection and Response	Network detection and response involves the collection and analysis of network event data to aid in the detection of malicious activity.	Agencies may consider network detection and response protections in cloud environments. These solutions can be cloud-native, third-party provided or, where possible, directly deployed and managed by the agency. Agencies may also consider external mechanisms like CASBs or other Security-as-a-Service capabilities to enable network detection and response. These could take the form of network detection and response as managed services, or services that provide network traffic to agency-managed network detection and response capabilities.

4.4.2.7 Enterprise PEP Security Capabilities

Agencies need to understand how to integrate their cloud environments with their existing environments to enable holistic, enterprise-wide accessibility, visibility, and management. As cloud environments are commonly accessible to external entities, often over the internet, agencies need to take extra care to secure their cloud resources and the mechanisms used to access them. Since new resources can be easily deployed in cloud environments, agencies must employ the necessary visibility and monitoring capabilities to understand which resources are being deployed and how users are accessing these resources. Table 9 lists the applicable Enterprise PEP Security Capabilities for the IaaS, PaaS, and SaaS guidance in the Cloud Use Case.

Table 9: Enterprise PEP Security Capabilities for IaaS, PaaS, and SaaS

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
Security Orchestration, Automation, and Response	Security Orchestration, Automation, and Response (SOAR) tools define, prioritize, and automate the response to security incidents.	When selecting cloud services, agencies should consider how the service would be integrated into their existing SOAR solution, including how telemetry may be made available and how automated responses might be handled. Agencies may consider using SOAR solutions that are natively available in each cloud environment but need to understand orchestration and response holistically across all their environments. If agencies are considering transitioning from an on-premises to a cloud-based SOAR solution, they should account for how the loss of connectivity by the agency campus or branch offices might affect automated responses.
Shadow Information Technology Detection	Shadow IT detection systems detect the presence of unauthorized software and systems in use by an agency.	Agencies should employ systems to detect the use of unsanctioned CSPs and the use of unsanctioned services in sanctioned CSPs. Agencies should consider managing cloud environments using methods that enable the automatic detection and potential remediation of unauthorized or noncompliant deployments in those environments. Agencies may consider updating and retraining users on workflows and administrative controls for subscribing to new services for official use.
Virtual Private Network	VPN solutions provide a secure communications mechanism between networks that may traverse across unprotected or public networks.	<p>Agencies need to ensure that only secure protocols are available for use for VPNs. VPN connections should require strong authentication, including the use of MFA for users connecting directly to cloud environments. Additionally, agencies should consider integrating endpoint compliance checking and remediation as part of authorizing VPN access. Agencies that use site-to-site VPNs to connect agency campuses to cloud environments should ensure that only authorized users and services are permitted to traverse the established VPN.</p> <p>Agencies should employ techniques to limit the access of entities connecting via VPN (e.g., network segmentation, application gateways, virtual desktop infrastructure (VDI)). Agencies should consider isolating the cloud VPN resources from their other cloud resources, and should also consider mechanisms that can more tightly align access with user, resource, or application workflows.</p> <p>Agencies need to ensure that VPN entry points are well-secured, including being up to date with security patches. Agencies need to consider the resiliency of these entry points to account for potential variations in authorized use and opportunities for denial of service. For VPN connections between agency campuses and cloud environments, agencies should consider backup paths with automatic failover.</p>
Application Container	An application container is a virtualization approach in which	Agencies may consider the use of application containers as part of their overall development and deployment infrastructure. Agencies need to ensure the security for any components that underly their application

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
	applications are isolated to a known set of dependencies, access methods, and interfaces.	containers, including performing appropriate configuration and patch management as appropriate. Agencies should consider protections to detect anomalous or malicious activities or behavior from application containers to mitigate the effects of container compromise, whether from direct attack on the application container or from supply chain compromise for the components that underly the application container. Given the brief nature of application container deployments, agencies will need to ensure that there is sufficient telemetry and logging available to enable effective incident response.
Remote Desktop Access	Remote desktop access solutions provide a mechanism for connecting to and controlling a remote physical or virtual computer.	<p>Agencies may consider enabling a cloud-hosted desktop using an agency-managed VDI or vendor-managed DaaS, depending on agency risk tolerance. When selecting cloud-hosted desktop solutions, agencies with existing VDI infrastructure may consider opportunities for integration between their existing infrastructure and cloud-hosted solutions. Agencies need to ensure resources are available to enable cloud-hosted desktop solutions to handle variations in authorized use. Agencies need to ensure that desktop instances are well-secured, including being up to date with security patches. Agencies may consider the use of ephemeral desktop instances to update desktop instance configurations and limit the persistence of compromised instances.</p> <p>Remote desktop access may be provided as a direct service or in combination with a VPN. Remote desktop access should only be made available using secure protocols and with strong authentication, including MFA and (possibly) endpoint compliance checking. Agencies may consider treating the desktop instances accessed via RDA like remote user devices and thus align the security protections in the cloud deployment to limit the trust placed in the desktop instances. Agencies may also consider protections like gateways or bastion hosts to limit direct access to desktop instances. For example, agencies should consider employing a multi-tier architecture that allows a front-end tier for user authentication and authorization, thereby applying contextual security filters based on user or device location, operating system, and other factors. Agencies should consider preventing local file saving and peripheral use, as well as strict enforcement of access application security settings.</p>
Costs Monitoring	Costs monitoring entails the monitoring of costs incurred by enterprise resources.	Cloud environments often make it easy to deploy new cloud resources or for mechanisms like elastic expansion to automatically expand the footprint of existing cloud resources. These features can facilitate more agile workflows which can respond faster to user needs. However, without appropriate constraints on their use, these features can be used to increase in the costs incurred by the agency, including both accidental misconfiguration as well as attacks that increase utilization through elastic expansion. Agencies should monitor their cloud resource usage to detect irregular resource usage.

4.4.2.8 Data Protection PEP Security Capabilities

Data protection is the process of maintaining the confidentiality, integrity, and availability of an agency's data consistent with their risk management strategy. It is important that agencies secure their data from corruption, compromise, and loss. Agencies should have processes and tools in place to protect agency data, prevent data exfiltration, and ensure the privacy and integrity of data, considering that data may be accessed from devices beyond the protections and perhaps administration of agencies. Agencies do not have control over physical protections for data stored at CSPs. Therefore, the application of data protection security capabilities is critical to securing agency data in all cloud deployments. Agencies should consider the sensitivity of data when

applying rigor to these Data Protection PEP Security Capabilities. Policies, procedures, and incident response may require adaptations to accommodate cloud storage and uses.

Table 10 lists the applicable Data Protection PEP Security Capabilities for the IaaS, PaaS, and SaaS guidance in the Cloud Use Case.

Table 10: Data Protection PEP Security Capabilities for IaaS, PaaS, and SaaS

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
Access Control	Access control technologies allow an agency to define policies concerning the allowable activities of users and entities to data and resources.	<p>To protect agency data in cloud environments, agencies should utilize fine-grained access controls that can account for context around data accesses (e.g., attribute based access controls (ABAC), Conditional role based access controls (RBAC)). When implementing these access control for cloud resources, it is important that agencies account for service, identity, device, tags (see Data Labeling, next row) and any other policies that are applied to the type of data.</p> <ul style="list-style-type: none"> • SaaS: Most SaaS applications have access control capabilities built in. Prior to migrating agency data to SaaS deployments, agencies should determine if these meet their needs and risk tolerance level. • PaaS: Some PaaS computing platforms (e.g., databases) will have native capabilities for access control. Agencies should determine if these meet their needs and risk tolerance level.
Protections for Data at Rest	Data protection at rest aims to secure data stored on any device or storage medium.	In line with Executive Order 14028, ³⁰ agencies should ensure that any agency data stored in a cloud environment is encrypted to mitigate the risk of loss and to ensure that data is inaccessible in the event of a breach of the agency cloud deployment or the CSP. Additionally, agencies need to consider the storage, access, and lifecycle of the keys used to encrypt or decrypt agency data to limit the opportunities for unauthorized access to the data. Data access and security protections should be in alignment with agency risk determinations and data sensitivity and, where possible, with policies in place to automatically align the data access and security protections in response to changes in risk determination or data sensitivity.
Protections for Data in Transit	Data protection in transit, or data in motion, aims to secure data that is actively moving from one location to another, such as across the internet or through a private enterprise network.	With agency users at the agency campus, branch offices, or working remotely will access agency data at a CSP, protections for data in transit are paramount. Beyond the encryption required by Executive Order 14028, ³¹ these data protections should combine identity guarantees of the recipient and validation of the endpoint receiving the data. Agencies must also consider how agency data is moved to a cloud deployment during the initial migration.

³⁰ “Executive Order 14028 Improving the Nation’s Cybersecurity,” Office of Management and Budget (May 2021). <https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/>.

³¹ “Executive Order 14028 Improving the Nation’s Cybersecurity,” Office of Management and Budget (May 2021). <https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/>.

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
Data Loss Prevention	DLP technologies detect instances of the exfiltration, either malicious or accidental, of agency data.	<p>Agencies should, where possible, consider data loss prevention in a holistic manner throughout the cloud environment and across the agency enterprise, ensuring data loss can be detected and potentially prevented when occurring over multiple modes of conveyance. Agency data should be brought into or out of cloud deployments using authentication and transport mechanisms that can protect the confidentiality and integrity of the data in line with agency risk tolerances. DLP solutions at agency cloud deployments should ensure that sensitive data does not enter cloud deployments unless it is authorized to be stored in the cloud deployment.</p> <p>DLP solutions at agency cloud deployments must also consider any possibilities for file exfiltration from the cloud. CSPs may offer DLP services that can be added to a service offering. Alternatively, agencies can implement DLP controls.</p> <ul style="list-style-type: none"> • SaaS: SaaS may directly integrate DLP into the service or may support the integration of third-party DLP controls. Agencies must understand what is provided by these offerings and if these solutions are robust enough to detect sensitive agency data. • PaaS/laaS: DLP may be implemented as part of egress controls or integrated in controls that apply throughout the deployment.
Data Access and Use Telemetry	Data access and use telemetry identifies agency-sensitive data stored, processed, or transmitted, including those located at a service provider, and enforces detailed logging for access or changes to sensitive data.	An agency should track access to all agency data and applications in the cloud and analyze all access events for suspicious behaviors. Most CSPs have native capabilities for logging, monitoring, and analysis of data access and use telemetry. Data access logs for cloud deployments can become quite large, so agencies may consider configuring data access logs to meet their individual needs.
Data Labeling	Data labeling is the process of tagging data by categories to protect and control the use of data and identify the risk level associated with the data.	<p>Agencies should deploy security tags with data in agency cloud deployments so that the tag can inform data access and security. If available, agencies should use automated data labeling through content inspection and ML. Most CSPs have the concept of a label or tag. It is important to have a list of tags that can be used consistently across multiple CSPs and cloud deployments. Some CSPs have rules that provide automation so that untagged or incorrectly tagged data can be identified.</p> <ul style="list-style-type: none"> • SaaS: Some SaaS deployments (e.g., business applications, collaborative workspaces) may have native capabilities that support data tagging and/or labeling. • PaaS: Some PaaS computing platforms (e.g., databases) may have native capabilities for automated data tagging.

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
Data Inventory	Data inventory entails developing, documenting, and maintaining a current inventory of agency data.	As agencies move data and applications into the cloud, they should maintain an inventory of data assets, which includes data assets stored in cloud deployments in the agency's inventory of datasets. The data inventory should include relevant metadata so agencies can securely locate, manage, and use data. Agencies should update this inventory as necessary when cloud deployments are created, modified, and retired. Most CSPs provide tools to inventory storage locations. Agencies should come up with enterprise solutions for data inventories when multiple CSPs are employed.

4.4.2.9 Services PEP Security Capabilities

The Services PEP capabilities apply when agency services in cloud environments are made available to users and entities outside the cloud environment, including agency users, external partners, other tenants in the same cloud environment or to external entities. Agency services are often a substantial target for threat actors. These services may contain access to agency data, which can be useful for attacking agency users and external entities. Agency data can also serve as an entry point for access into the agency enterprise network. When moving or deploying services into cloud environments, agencies need to understand the risks associated with the services, including their potential for misuse. Additionally, agencies need to understand the visibility they have into the overall environment that the application is deployed into, as well as the security controls that can be deployed to secure the applications.

Agencies may consider deployment models, like CDNs or CASBs, that deploy security protections in a different cloud environment or at a different vantage point in the cloud environment. To ensure all security protections are applied, agencies need to understand all the potential ways in which traffic may be received by the service to ensure that it only receives traffic sent through the security protections.

As in-line traffic decryption can affect end-to-end authentication, agencies will need to ensure that their services have strong end-to-end authentication in line with their risk tolerances. Additionally, agencies need to ensure that the services and networks with visibility into the decrypted traffic align with agency requirements and risk tolerances.

Beyond direct attacks against agency services, agencies should account for the potential compromise of the service prior to or during the deployment process. Agencies should consider approaches to development and deployment that consider security across the entire lifecycle, with verification and validation at each step in the process. Table 11 lists the applicable Services PEP Security Capabilities for the IaaS, PaaS, and SaaS guidance in the Cloud Use Case.

This capability group and all capabilities in Table 11 are new and will be added to the next version of the Security Capabilities Catalog during the next revision cycle.

Table 11: Services PEP Security Capabilities for IaaS, PaaS, and SaaS

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
Active Content Mitigation	Active content mitigation protections detect the presence of unapproved active content and facilitate its removal.	Agencies will need to ensure that active content mitigation protections align with the needs of and threats against their deployed services. Active content mitigation may need to be applied to all content being ingested, whether in application-level traffic or files. Agencies need to understand the protections employed, especially when they have limited visibility into traffic to the agency service. Agencies should consider applying mitigations to the content sent by the agency service. This will prevent the agency service from being used to attack agency or external users.
Data Loss Prevention	DLP technologies detect instances of the exfiltration, either malicious or accidental, of agency data.	Agencies should ensure that agency services only make data available using authentication and transport mechanisms that are able to protect the confidentiality and integrity of the data in line with agency risk tolerances. Agencies should consider how best to integrate agency service DLP as a component in their enterprise DLP solution to ensure that agency services are able to properly prevent data loss, and that data loss can be detected and potentially prevented when occurring over multiple modes of conveyance.
Protocol Compliance Enforcement	Protocol compliance enforcement technologies ensure that traffic complies with protocol definitions, like those documented by the IETF. ³²	Agencies should consider the use of proxies or other mechanisms for enforcing protocol compliance to help mitigate against the limited visibility that agencies may have into low-level-network traffic details in cloud environments. Agencies will need to understand all the potential ways the service can be accessed to ensure that all accesses are mediated by these compliance mechanisms.
Malicious Content Filtering	Malicious content filtering protections detect the presence of malicious content and facilitate its removal.	Agencies may need to tune the malicious content filtering to account for specific threats that may apply to a given agency service. Malicious content filtering may need to be applied to all content being ingested into the agency, whether in application-level traffic or files. Agencies need to understand the protections employed, especially when they have limited visibility into traffic to the agency service. Agencies should consider malicious content filtering for content sent by the agency service to prevent agency services from being used to attack agency or external users.
Access Control	Access control technologies allow an agency to define policies limiting what actions may be performed by connected users and entities.	Agencies should understand how to integrate agency service access controls into their overall enterprise entitlement management workflow to ensure visibility into and control over all resources that entities have access. Agencies should consider enabling MFA to access agency services to mitigate the effects of password compromise, device loss or theft, or device impersonation. Agencies should consider authentication strength according to user roles, device security posture and compliance, and anomalous or suspicious login or user behavior.

4.4.2.10 Identity PEP Security Capabilities

As cloud resources are made available to agency users and external entities, identity forms a key component in protecting those assets. This need for strong identity protections is especially important as agency users often access these resources from remote locations, which can provide a more limited view into user devices and environments. Given the opportunities for account compromise and the potential for exfiltration or lateral movement from the cloud environment, agencies need to consider protections beyond simply identity

³² “RFCs,” Internet Engineering Task Force (2021). <https://www.ietf.org/standards/rfcs/>.

authentication. Some of these protections including device security posture and compliance checking, as well as anomalous or suspicious user behavior detection.

While agencies can have enterprise identity stores for their users, cloud environments often contain their identity stores, with agency users potentially having identities in both locations. Agencies may consider ways for integrating these environments, including single sign-on and federated identities. Agencies may also consider moving the identity store out of their traditional on-premises environment and into a cloud environment. For each of these models, the agency needs to consider the potential for lateral movement through their environments by compromised accounts or forged identities. Additionally, as identities become more distributed, agencies need to ensure they retain enterprise-wide visibility into their identities and the accesses of those identities. Table 12 lists the applicable Identity PEP Security Capabilities for the IaaS, PaaS, and SaaS guidance in the Cloud Use Case.

This capability group and all capabilities in Table 12 are new and will be added to the next version of the Security Capabilities Catalog during the next revision cycle.

Table 12: Identity PEP Security Capabilities for IaaS, PaaS, and SaaS

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
Adaptive Authentication	Adaptive authentication aligns the strength of the user or entity authentication mechanisms to the level of risk associated with the requested authorization.	Agencies should consider authentication strength according to user roles, device security posture and compliance, and anomalous or suspicious login or user behavior.
Entitlement Inventory	Entitlement inventory entails developing, documenting, and maintaining a current inventory of user and entity permissions and authorizations to agency resources.	Agencies may use cloud-native tools to maintain per-cloud entitlement inventories but should consider methods for integrating these per-cloud inventories to provide visibility across all agency environments.
Service Identity	Service identity ensures that users and entities can authenticate the identities of agency services.	Agencies should ensure that all services available from outside the cloud environment have an appropriate identity that allows for mutual authentication. The identities should, when possible, be securely tracked and managed following lifecycle policies appropriate to the security of the service. Agencies should consider allowing for mutual authentication for all internal and external services and should consider enabling mutual authentication for all data flows.
Secrets Management	Secrets management entails developing and using a formal process to securely track and manage digital authentication credentials, including certificates, passwords, and API keys.	Agencies should consider managing cloud secrets using secrets management systems that facilitate lifecycle management and secure storage and access. Agencies may consider methods to manage secrets in a unified manner across all agency environments. Agencies should ensure that secrets management is strongly integrated into their development and configuration management processes.

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
Behavioral Baselineing	Behavioral baselining is capturing information about user and entity behavior to enable dynamic threat discovery and facilitate vulnerability management.	Agencies should understand and account for user behavior in their cloud environments to allow for the detection of anomalous or malicious behavior. When obtaining user and entity baselines, agencies should consider accounting for both activities performed in the cloud environment as well as information about where and how users or entities access cloud resources. Agencies may deploy baselining solutions in the cloud environment or may integrate user and entity activity into an externally deployed solution. Agencies should account for how to integrate cloud-specific behavioral baselines into an overall baseline to understand user and entity behavior holistically across the enterprise.
Enterprise Identity, Credential, and Access Management	Enterprise ICAM entails maintaining visibility into agency identities across agency environments and managing changes to those identities through a formal (preferably automated) process.	Agencies integrating ICAM to provide a common user identity should work to minimize opportunities for lateral movement across the environments. This includes strong application of least privilege, limiting privileged accounts, and enabling detections for anomalous or malicious user and entity behavior. Agencies should consider methods for ensuring continuous compliance of permissions and identities across cloud environments. Agencies may use CIEM tools to facilitate the enterprise-wide management of least privilege. ³³
Multi-factor Authentication	MFA entails using two or more factors to verify user or entity identity.	Agencies should, wherever possible, employ phishing-resistant MFA as part of verifying identity. ³⁴ Agencies should consider re-verification of identity when users or entities seek to perform suspicious or sensitive actions. This will allow agencies to minimize opportunities for lateral movement or privilege escalation from compromised machines or devices. When protecting a cloud resource with MFA, agencies need to understand all methods for accessing that resource to ensure there are no alternative routes that can bypass MFA.
Continuous Authentication	Continuous authentication entails validating and re-authenticating identity through the lifecycle of entity interactions.	Agencies should employ solutions that re-verify identity when users or entities seek to perform sensitive actions or when anomalous or suspicious behavior is detected. This includes aligning the strength of authentication for the re-verification according to user roles, device security posture and compliance, and the sensitivity of the requested action.

³³ "Enabling Mission Delivery through Improved Identity, Credential, and Access Management," Office of Management and Budget (May 2019). <https://www.whitehouse.gov/wp-content/uploads/2019/05/M-19-17.pdf>.

³⁴ "Moving the U.S. Government Toward Zero Trust Cybersecurity Principles," Office of Management and Budget M-22-09 (2022). <https://www.whitehouse.gov/wp-content/uploads/2022/01/M-22-09.pdf>.

4.5 TELEMETRY REQUIREMENTS

As agencies transition from on-premises deployments to deployments in cloud environments, visibility by CISA must be preserved through information sharing. Figure 13 shows the conceptual architecture of the IaaS, PaaS, and SaaS guidance in the Cloud Use Case, with the telemetry requirements added as lines on certain data flows. These lines, depicted in Figure 13, indicate when an agency must share telemetry with CISA. Subject to applicable law, CISA may require internal telemetry to be collected. The requirements for sharing telemetry data with CISA are only applicable to the data flows between the cloud environment, web, external partners and external entities. It is assumed that Agency and Remote Users are considered internal -agency traffic. Consult the NCPS program³⁵ and CDM program³⁶ for further details.

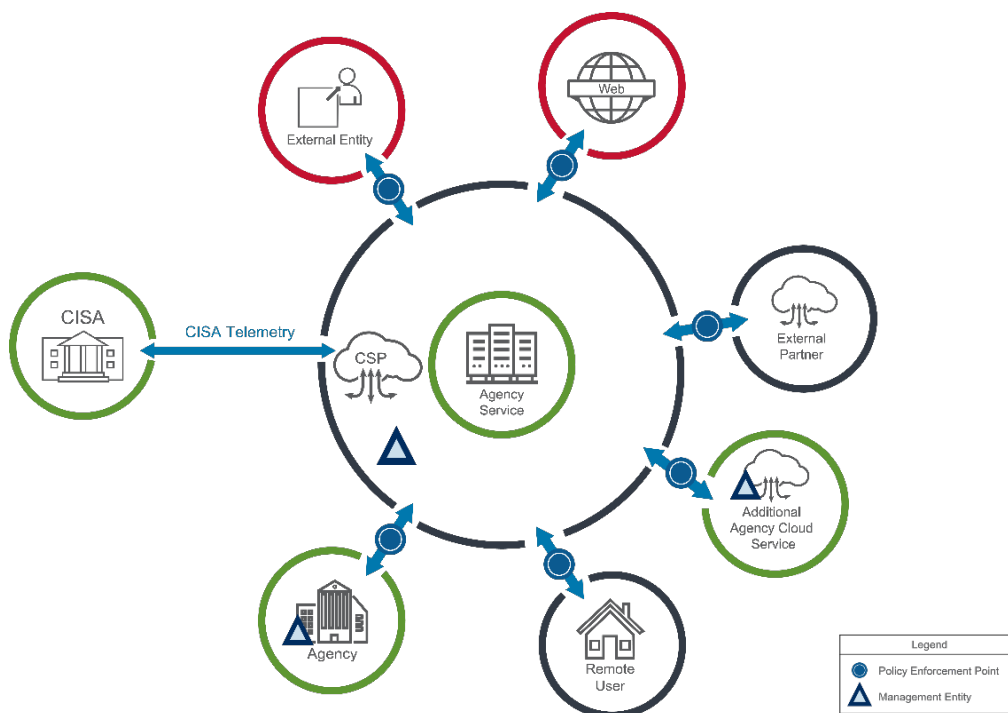


Figure 13: IaaS, PaaS, and SaaS Telemetry Sharing with CISA

³⁵ "National Cybersecurity Protection System," Cybersecurity and Infrastructure Security Agency. <https://cisa.gov/national-cybersecurity-protection-system-ncps>.

³⁶ "Continuous Diagnostics and Mitigation," Cybersecurity and Infrastructure Security Agency. <https://cisa.gov/cdm>.

5. EMAIL-AS-A-SERVICE USE CASE

This section broadly covers Email-as-a-Service (EaaS) deployment models, as outlined in OMB M-19-26. It defines how network and multi-boundary security should be applied when an agency's email service is hosted by a cloud service provider and the provider is responsible for the email infrastructure.

This section builds upon the guidance on IaaS, PaaS, and SaaS deployment models, as outlined in Section 4. Agencies should refer to Section 4 for general cloud and SaaS guidance.

This section includes three network security patterns:

- Agency campus users accessing email or sending email through the agency-sanctioned EaaS provider.
- Agency remote user accessing email or sending email through the agency-sanctioned EaaS provider.
- External entity sending email to or receiving email from agency-sanctioned EaaS provider.

An agency may implement a subset of these security patterns (and not necessarily all three), depending on how agencies are migrating and deploying services in the cloud.

Agencies may implement **additional security patterns** not covered in the EaaS Use Case.

Agencies may implement additional security patterns not covered in the EaaS Use Case. These additional security patterns may be in scope for a different use case but would be out of scope of the EaaS guidance in the Cloud Use Case.

5.1 ASSUMPTIONS AND CONSTRAINTS

This section outlines guiding assumptions and constraints for the EaaS guidance in the Cloud Use Case. It is intended to clarify significant details about the construction and replication of the EaaS guidance in this use case. The assumptions are broken down by the EaaS guidance in this use case as a whole and by the unique entities discussed in this section:

- Agency campus
- Agency EaaS provider
- Remote users
- External entities

The following are the assumptions and constraints of the EaaS guidance in this use case:

- Requirements for information sharing with CISA in support of NCPS and CDM purposes are beyond the scope of this document. Consult the NCPS³⁷ program and CDM³⁸ program for further details.
- Requirements for endpoint protection are beyond the scope of this document. Consult the FISMA or NIST references in Appendix B for additional guidance on endpoint protections, BYOD, and telework security.

³⁷ "National Cybersecurity Protection System," Cybersecurity and Infrastructure Security Agency. <https://cisa.gov/national-cybersecurity-protection-system-ncps>.

³⁸ "Continuous Diagnostics and Mitigation," Cybersecurity and Infrastructure Security Agency. <https://cisa.gov/cdm>.

- The TIC security capabilities applicable to the use case do not depend on a particular data transfer mechanism. In other words, the same capabilities apply if the conveyance is over leased lines, software VPN, hardware VPN, etc.
- The EaaS guidance in the Cloud Use Case is primarily focused on network security. While this use case can be compatible with zero trust, implementation of zero trust requires additional controls and measures beyond those detailed in this use case.

The following are assumptions about the agency campus:

- For this use case, the agency campus may refer to an agency's main campus, branch office, or both.
- The agency campus utilizes the Traditional TIC Use Case, or equivalent security architectures, to access the web and CSPs.
- Any branch office may utilize the Branch Office Use Case, or equivalent security architectures, to access the web, CSPs, and the agency campus.
- The agency maintains control over and has significant visibility into the agency campus.
- Data is protected at a level commensurate with the agency's risk tolerance and in accordance with federal requirements.
- The agency employs NOC and SOC tools capable of maintaining and protecting their portions of the overall infrastructure. To accomplish this, agencies can opt to use an NOC and SOC, or commensurate solutions.

The following are assumptions about agency-sanctioned EaaS providers:

- EaaS providers are compliant with FedRAMP.³⁹
- Interactions with service providers follow agency-defined policies and procedures for business need justification, partner connection eligibility, service levels, data protections, incident response information sharing and reporting, costs, data ownership, and contracting.
- The agency maintains awareness of the email providers that are sanctioned for use by the agency. The agency may use this awareness to limit access to certain email services on approved providers.
- The agency has limited control over and visibility into EaaS provider environments relative to other entities, like the agency campus.
- All agency-generated email is sent to external entities through one or more agency-sanctioned EaaS providers.
- All email from external entities to the agency is received through and stored on one or more agency-sanctioned EaaS providers.
- EaaS providers have NOCs and SOCs that control and protect the portions of the service infrastructure where the agency has little or no control or visibility.
- The agency only uses secure mechanisms (e.g., TLS, VPN etc.) for EaaS administration.
- The agency only uses strong authentication mechanisms (e.g., FIPS 140-3⁴⁰ compliant MFA for EaaS administration).
- Data stored at EaaS providers is protected at a level commensurate with the agency's risk tolerance and in accordance with federal requirements.
- EaaS providers allow the agency to define and/or configure policies that the provider applies on their behalf.
- EaaS providers allow the agency to define roles and responsibilities for the definition and configuration of policies applied on their behalf by the provider.
- EaaS providers have mechanisms that allow the agency to obtain visibility into the current state and history of the system (e.g., log information, configuration, accesses, system activity).

³⁹ "FedRAMP," General Services Administration (2019). <https://www.fedramp.gov/federal-agencies/>.

⁴⁰ "FIPS 140-3 NIST Security Requirements for Cryptographic Modules," National Institute of Standards and Technology (2019). <https://csrc.nist.gov/publications/detail/fips/140/3/final>.

- EaaS providers enable commensurate protections and policy enforcement for traffic between the agency tenant and other tenants of the provider as between the agency tenant and parties outside the provider.

The following are assumptions about remote users:

- The remote user utilizes the Remote User Use Case, or equivalent security architectures, to access the agency campus, the web, and CSPs.
- The remote user may be using either GFE or BYOD.
- For GFE, remote users may be permitted business-only use of their devices (e.g., COBE) or permitted for personal use (e.g., COPE).
- Devices employed by remote users may include desktops, laptops, and mobile devices (e.g., smartphones, tablets). While remote users may connect to virtual desktop instances hosted by the agency or in cloud service providers, these agency-managed desktop instances are not considered remote user devices. However, they may be considered as agency virtual GFEs inside an agency campus or cloud environment.
- Agencies may have limited control over or visibility into devices used by the remote user.
- Email traffic to and from the remote user devices is in scope for the EaaS guidance of the Cloud Use Case. Other traffic may be in scope for other use cases.
- Agency data on remote user devices, or in transit to and from them, is protected at a level commensurate with the agency's risk tolerance and in accordance with federal requirements.
- The agency employs NOC and SOC tools capable of protecting remote user sessions. These functions may be performed as an extension to the NOC and SOC tools managed and housed at the agency campus or via commensurate solutions.

The following are assumptions about external entities:

- External entities include public users sending and receiving email to and from agency email service.
- The agency may not be able to rely on policies deployed by external entities.

5.2 CONCEPTUAL ARCHITECTURE

The EaaS guidance in the Cloud Use Case focuses on the scenario in which an agency is using a cloud deployment for its agency email.

As shown in Figure 14, this conceptual architecture is composed of four distinct trust zones: agency campus, EaaS provider, remote user, and external entity. This conceptual architecture shows a single remote user and a single external entity trust zone. These simplifications are not meant to imply that an agency must treat all remote users or external entities in the same manner. Applicable TIC capabilities and their rigor should be tailored for the nature of the remote user, external entity, or EaaS provider.




Figure 14: EaaS Conceptual Architecture

The trust zones depicted in Figure 14 are detailed in Table 13. The trust zones are labeled with levels of trust, using the example trust levels—high, medium, and low—explained in the Reference Architecture. While the trust levels assigned to each of these zones were selected based on existing pilots or deployments, the trust assignments may not capture the needs or requirements of all agencies. Agencies may assign different trust levels to trust zones, based on their own risk tolerance. For example, an agency might decide to designate a EaaS provider with a higher trust level based on agency criteria (e.g., the accreditation level of the EaaS provider, the control and visibility, available protections). Additionally, an agency may have remote users that employ unmanaged personal devices and may decide to label remote users with a lower trust level.

Implementation Consideration

The trust levels in this use case are intended to be examples. Agencies may define and assign trust levels to align with their requirements, environments, and risk tolerance.

Table 13 briefly explains why each entity is labeled with either a high, medium, or low trust zone level in this conceptual architecture to help agencies determine what is most appropriate in their implementation.

Table 13: Trust Zones in the Cloud Use Case for EaaS

Trust Zone	Description
Agency Campus Trust Zone	The Agency Campus Trust Zone is the logical zone for the agency campus or the agency's enterprise network. The trust zone includes MGMTs such as the NOC, SOC, and other entities. The agency maintains control over and visibility into the agency campus. The agency campus employs the Traditional TIC or Branch Office Use Cases, or equivalent, including when transmitting traffic from the EaaS provider to external entities. The Agency Campus Trust Zone is labeled with a <i>high trust level</i> in this use case.
EaaS Provider Trust Zone	The EaaS Provider Trust Zone is a logical trust zone for the CSP providing email service to the agency. EaaS deployments follow a shared responsibility model, with the EaaS provider

Trust Zone	Description
	responsible for protecting the underlying cloud infrastructure and the agency providing certain policy-defined functions and capabilities. The trust zone includes a MGMT that executes locally scoped functions for the EaaS environment. The EaaS Provider Trust Zone is labeled with a medium trust level in this conceptual architecture due to the potential for limited agency control over and visibility into the EaaS environment.
Remote User Trust Zone	The Remote User Trust Zone is a logical trust zone representing a device employed by a remote user when accessing the EaaS provider. Remote user devices may be agency-managed (e.g., GFE) or not managed by agencies (e.g., BYOD). Devices not managed by agencies may not be suitable for performing some policy enforcement capabilities. The agency may have no control over or visibility into non-GFE devices and may have limited control over or visibility into agency-managed devices. The remote user employs the Remote User Use Case. The Remote User Trust Zone is labeled with a medium trust level in this conceptual architecture.
External Entity Trust Zone	The External Entity Trust Zone is a logical zone that depicts an unmanaged, and potentially untrusted, external entity communicating with agency entities through sending or receiving email via the agency email service, and with no PEPs or MGMTs where the agency, or entities acting on its behalf, may deploy policies. An external entity may depict a nonhuman entity (e.g., an email service). Given these limitations, the External Entity Trust Zone is labeled with a low trust level in this conceptual architecture.

5.2.1 Risk and Deployment Considerations

As agencies migrate their corporate email from on-premises deployments to cloud deployments, they must understand the differences between the two models, how to protect the cloud deployment, how the agency security posture must adapt, and best practices for mitigating inherent risks. Email is a critical application of all agency users and includes internal agency communications, communication with external partners, and communication with the public. Thus, email contains agency-sensitive information and agencies must consider their risks and ensure that there are email security controls, policies, and operational processes in place to manage risks. Additionally, email is likely an agency's largest attack surface and, consequently, is a significant attack vector for malware and credential theft.

Many risk considerations inherent in a cloud deployment were discussed in the IaaS, PaaS, and SaaS guidance of this use case. These will not be repeated here, so agencies should refer to Section 4.2.2 when making risk-informed decisions about their email solution in the cloud.

5.2.1.1 Email Attacks and Threats

Email is among the most common vectors used to attack agency networks. These email-based attacks provide threat actors with the initial access from which they can persist and move laterally throughout the enterprise. There are numerous potential email threats, including:

- **Spam:** Electronic junk mail or the abuse of electronic messaging systems to indiscriminately send unsolicited bulk messages.
- **Phishing:** A technique for attempting to compromise an account or to acquire sensitive data through a fraudulent solicitation in email, in which the perpetrator masquerades as a legitimate business or reputable person.
- **Spear Phishing:** A targeted phishing attack against a specific user or group.
- **Whaling:** A targeted phishing attack against high-ranking members of organizations.

- **Malicious Attachments:** An attachment to email that is designed to launch an attack on a computer or system, potentially with obfuscated attachment information to deceive users. These can include ransomware, spyware, malicious PDFs, documents, voicemails, and disguised files. Disguised files deceive users by having a benign filename to hide the true malicious behavior.
- **Malicious Links:** Links included in an email that direct users to malicious websites. These could include websites that attempt to compromise accounts or acquire sensitive data, or websites or files designed to launch an attack on the user's computer or system.
- **Spoofing:** Faking the sending address of a malicious email to increase the likelihood of the recipient taking the desired action.
- **Email Service Attacks:** Attacks directed at the email service itself. These could compromise the service or interfere with the sending or receiving of emails.

As agencies migrate from on-premises to EaaS deployments, they will need to understand the relevant email threats along with the native and third-party security capabilities available for the EaaS provider to ensure proper alignment.

5.2.1.2 Email as a Mission Critical Agency Application

Email is the most widely used tool for communication by agencies. It is used by agency employees and contractors daily. It is a core agency service and a key part of agency workflow. Since it is used extensively for communication with partners and the public, it is critical that agencies have security policies and procedures in place that explicitly consider email usage and transmission. Agencies should employ robust monitoring of email traffic and auditing of access to the agency email service. Agencies should use all email traffic and access logs for threat detection and discovery.

A complex attack to an agency's email service or the agency's EaaS provider can cause a significant disruption to the agency's operations. In addition to disrupting the agency's mission critical operations, it can also disrupt an agency's incident response procedures. Agencies should have an incident response plan that includes notifying the security team and users of a compromised email service without using email. The agency incident response team should avoid the use of email for any incident response activities so that the threat actor does not detect incident response activities.⁴¹

5.2.1.3 Potential for Access to Agency Data

Agency email represents a substantial fraction of the agency user communication, both internally and externally, including potentially sensitive or private information. The information available in these communications may have compliance rules regarding access, handling, and destruction. Traditional on-premises deployments of email solutions enabled agencies to align the protections, both physical and digital, according to sensitivity or compliance needs. EaaS is the shared responsibility model, with both the agency and the provider having access to the service. To provide assurances of commensurate protections, service providers often provide SLAs, potentially including audits by third-party vendors. As agencies commonly have limited direct visibility into the protections employed by the service provider, they often must rely on these assurances for alignment with their risk tolerances.

There may be opportunities for agencies to enable capabilities that limit the visibility of this data while stored in the email service (e.g., requiring all email sent or received be encrypted). These capabilities, however, can increase the complexity for end users or external entities and can limit the features available for use by the agency.

⁴¹ "Federal Government Cybersecurity Incident and Vulnerability Response Playbook," Cybersecurity and Infrastructure Agency (2021).
https://cisa.gov/sites/default/files/publications/Federal_Government_Cybersecurity_Incident_and_Vulnerability_Response_Playbooks_508C.pdf.

5.2.1.4 Potential Channel for Data Exfiltration

A primary use of agency email is to enable agency users and entities to communicate with external, potentially untrusted, entities. This communication method can provide a channel for threat actors to exfiltrate agency data. Agencies will need to understand the potential methods for exfiltrating data through cloud email deployments. For example, threat actors can use the remote accessibility offered to agency users to access the service using compromised accounts and retrieve data to attacker-controlled endpoints. Threat actors can also configure policies to automatically send newly received emails to attacker-controlled external accounts to continue exfiltration without needing to access the service. Beyond exfiltrating agency data available in existing emails, threat actors can use their access to mailboxes to store data that they have collected from the agency for eventual exfiltration. Compromised administrator accounts provide further opportunities to exfiltrate data, including methods like accessing backups or transferring data to other cloud tenants, whose visibility may not align with traditional exfiltration techniques. As agencies transition from on-premises email deployments to cloud deployments, they should understand the capabilities, visibility, and control available from their EaaS provider and align their detections, processes, and workflows to ensure their detections and protections for data exfiltration are in line with agency risk tolerances.

5.2.1.5 Third-Party Email Senders

Agencies commonly employ third-party services that send email on behalf of the agency (e.g., marketing, newsletters, mailing lists). These services often make use of their own infrastructure to send the email, sometimes routed through the agency email service. Traditional deployments provided opportunities for agencies to apply a variety of protections to ensure the security of these relays and that their services were not used to send unauthorized email. As agencies transition to EaaS, they may have fewer opportunities to deploy protections, increasing the opportunity for a misconfiguration to allow external entities to send unauthorized email.

Alternatively, some third-party DomainKeys Identified Mail (DKIM) services can send mail on behalf of the agency directly from the service provider's infrastructure without using the agency email service. In these scenarios, the third-party service needs to be authorized, using DKIM or Sender Policy Framework (SPF), to send mail on behalf of the agency.⁴² While this model can ease configuration and deployment, the agency may have limited visibility into email being sent on behalf of the agency, especially if the service is compromised or if the authorization to send mail is misconfigured.

5.2.2 Email-as-a-Service Connectivity

When selecting an EaaS provider, agencies will need to understand the options for connecting their campuses to the cloud environment and the options for enabling their users to access to the services. Section 4.2.3 of this use case more generally addresses connectivity for campuses and users to cloud environments. However, EaaS is often made available via direct connection from end users, so agencies may not have the same ability to use VPNs or remote desktop access to mediate access. While this access model can facilitate uniform security protections independent of access location, it offers the most opportunities for untrusted entities to potentially access the service, especially due to misconfiguration or account compromise.

5.3 SECURITY PATTERNS

Three security patterns capture the data flows for the EaaS guidance in the Cloud Use Case. Each of these has distinct sources, destinations, and options for policy enforcement. Regardless of the options chosen, agencies must ensure they are protecting their data in line with applicable federal requirements and agency risk tolerances, especially in instances where security policies are being applied by a third party on an agency's behalf, or in locations outside of the agency's traditional sphere of control. In cases where additional security

⁴² "Binding Operational Directive 18-01 Enhance Email and Web Security," Cybersecurity and Infrastructure Security Agency (2017). <https://www.cisa.gov/binding-operational-directive-18-01>.

capabilities are necessary to manage residual risk, agencies should apply the controls or explore options for compensating capabilities that achieve the desired protections to manage risks. The security patterns include the following trust zones:

- Agency campus
- EaaS provider
- Remote user
- External entity

The trust levels in these security patterns may not align with agency understanding of their environment; therefore, agencies may determine and label trust zones according to those that best describe their environment.

5.3.1 Security Pattern 1: Agency Campus User to Agency Email Service

Figure 15 illustrates the security pattern where users within the agency campus trust zone are accessing email resources. Two options are available for this connectivity and are outlined in Figure 15. Agencies may apply different constraints on connectivity options to different methods of accessing the agency email service. The EaaS provider may also impose requirements on connectivity. The agency should protect its data in line with its risk tolerances and federal requirements.

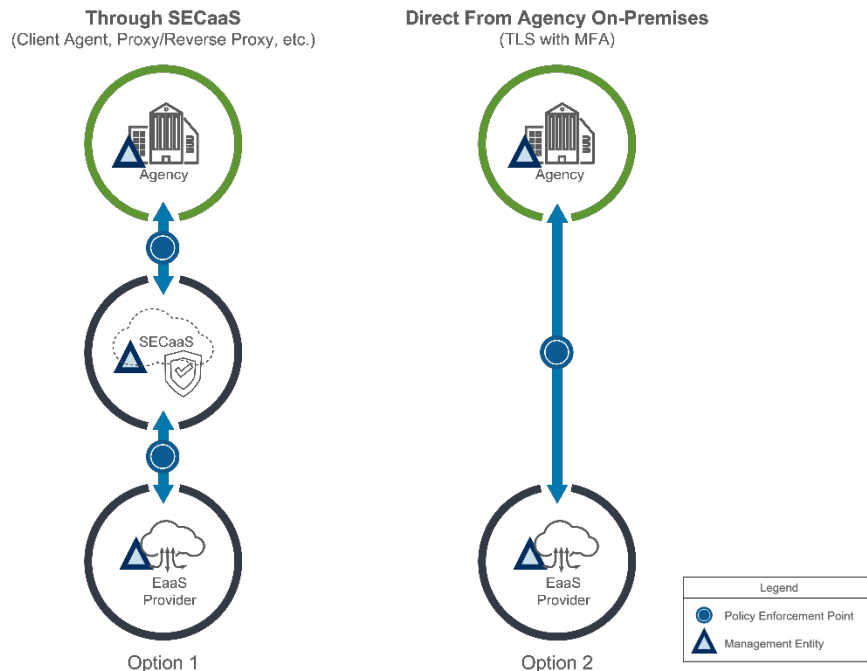


Figure 15: Security Pattern 1 – Agency Campus User to Agency Email Service

Through SECaaS
(Client Agent, Proxy/Reverse Proxy, etc)



Option 1

Direct From Agency On-Premises
(TLS with MFA)



Option 2

Option 1: The **first option** (left) permits connectivity from on-campus agency users to the agency email service via a SECaaS provider. Policy enforcement can be performed at the SECaaS provider, the agency campus, and the EaaS provider. Policy enforcement parity can be simplified when policy enforcement is handled at the SECaaS provider or EaaS provider. Various methods can be used to direct on-campus agency user traffic to the SECaaS provider, including client agents, proxy settings, transparent proxying, and DNS. The SECaaS trust zone is labeled with a medium trust level in this option, though agencies may determine and label trust zones according to the trust levels that best describe their environment.

Option 2: The **second option** (left) permits connectivity from on-campus agency users directly to the agency email service via protected connections (TLS with MFA, etc.). Policy enforcement can be performed at the agency campus and the EaaS provider. Policy enforcement parity across multiple campuses can be simplified when policy enforcement is performed at the EaaS provider.

5.3.2 Security Pattern 2: Agency Remote User to Agency Email Service

Figure 16 illustrates the security pattern where remote agency users are accessing the agency email service. Three options are available for this connectivity and are outlined in Figure 16. Agencies may apply different constraints on connectivity options to different methods of accessing the agency email service. The EaaS provider may also impose requirements on connectivity. An agency should protect its data in accordance with its risk tolerances and federal requirements.

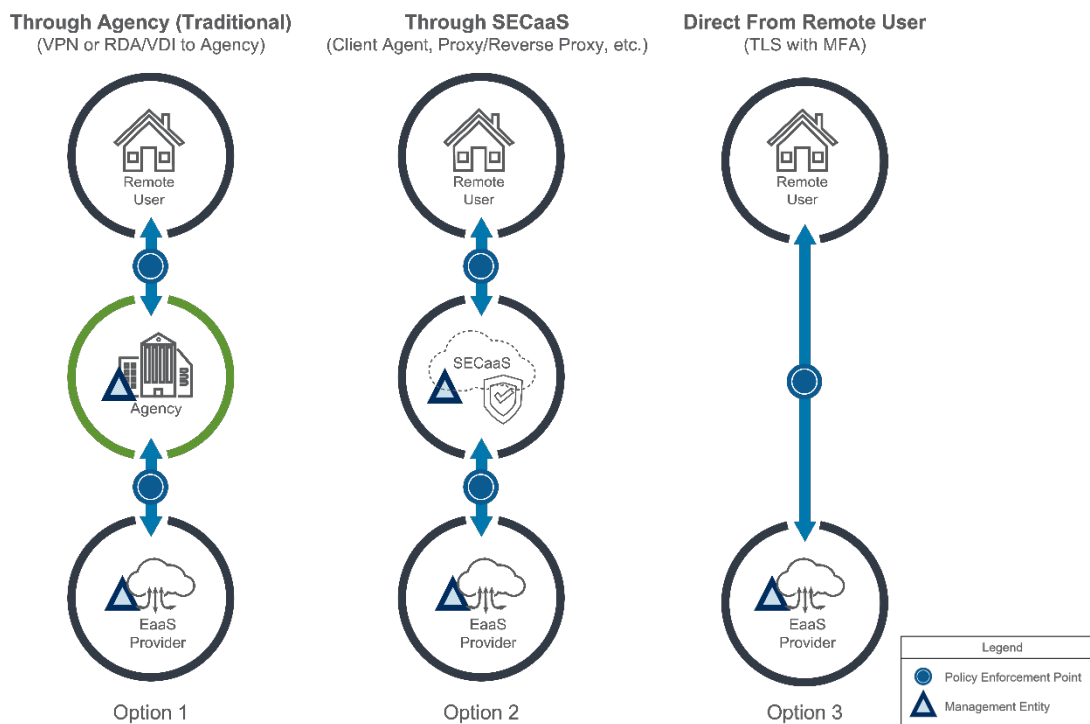


Figure 16: Security Pattern 2 – Agency Remote User to Agency Email Service

Through Agency (Traditional)
(VPN or RDA/ VDI to Agency)



Option 1: The first option (left) aligns with traditional mechanisms for remote users accessing agency email services. The remote user establishes a secure connection to the agency campus, as described in the Remote User Use Case, and accesses the email resources through that channel. Policy enforcement can be applied at the agency campus, the EaaS provider and, if possible, on the remote user's device. Policy enforcement parity between remote users and email resources can be simplified by applying protections at the agency campus or the EaaS provider.

Through SECaaS
(Client Agent, Proxy/Reverse Proxy, etc.)



Option 2

Direct From Remote User
(TLS with MFA)



Option 3

Option 2: The **second option** (left) permits connectivity from remote users to the agency email service via a SECaaS provider. Policy enforcement can be performed at the SECaaS provider, the EaaS provider and, if possible, on the remote user's device. Policy enforcement parity between cloud resources can be simplified when all cloud access passes through the same SECaaS provider. Various methods can be used to direct remote user traffic to the SECaaS provider, including client agents, proxy settings, transparent proxying, and DNS. The SECaaS trust zone is labeled with a medium trust level in this option, though agencies may determine and label trust zones according to the trust levels that best describe their environment.

Option 3: The **third option** (left) permits connectivity from remote users directly to the agency email service via protected connections (TLS with MFA, etc.). Policy enforcement can be performed at the EaaS provider and, if possible, on the remote user's device. Policy enforcement parity across users can be simplified when policy enforcement is performed at the EaaS provider.

5.3.3 Security Pattern 3: External Entity to Agency Email Service

Figure 17 illustrates the security pattern where an external entity communicates via email with the agency. Connections in this security pattern are among the riskiest as data is being received from or sent to potentially untrusted sources; therefore, a commensurate amount of rigor should be applied to the security capabilities. Three options are available for this connectivity and are outlined in Figure 17. Agencies may apply different constraints on connectivity options to different external entities. The EaaS provider may also impose requirements on connectivity. An agency should protect its data in accordance with its risk tolerances and federal requirements.

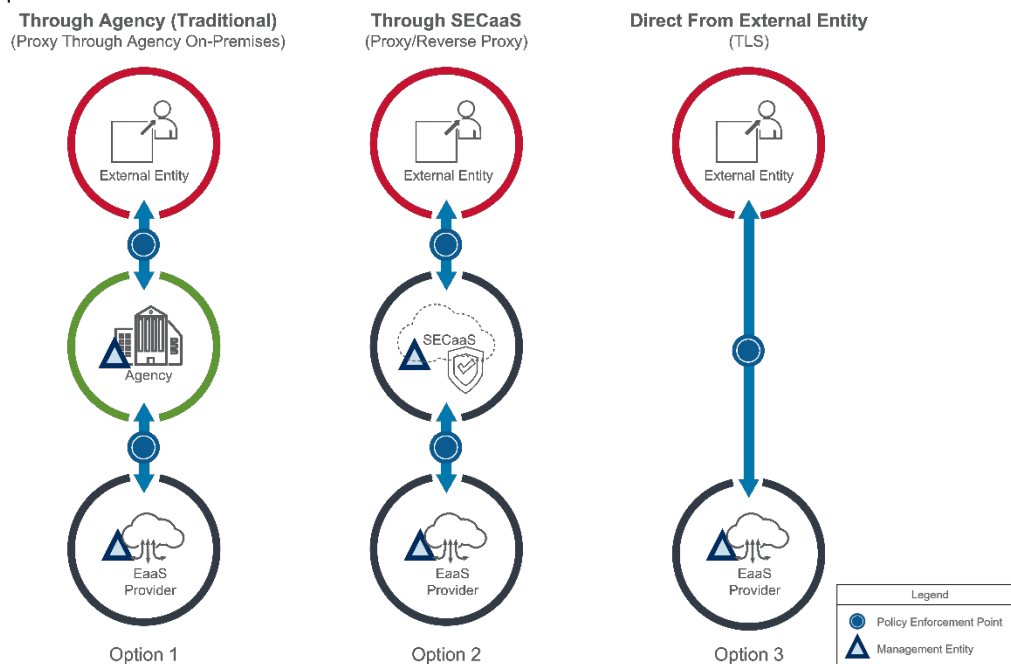



Figure 17: Security Pattern 3 – External Entity to Agency Email Service

Through Agency (Traditional)
(Proxy Through Agency On-Premises)



Option 1: The **first option** (left) aligns with traditional mechanisms for email traffic between external entities and agency email services. The agency campus acts as the front door to the agency email service, routing email between external entities and the service via secure connections. Policy enforcement can be applied at the agency campus and the agency email service.



Option 2: The **second option** (left) permits email traffic between external entities and the agency email service via a SECaaS provider. Policy enforcement can be performed at the SECaaS provider and the agency email service. Various methods can be used to direct external entity traffic to the SECaaS provider, including DNS and transparent proxying. The SECaaS trust zone is labeled with a medium trust level in this option, though agencies may determine and label trust zones according to the trust levels that best describe their environment.

Option 3: The **third option** (left) permits mail traffic directly between external entities and the agency email service, potentially via protected connections (TLS, etc.). Policy enforcement can be performed at the agency email service.

5.4 APPLICABLE SECURITY CAPABILITIES

The Security Capabilities Catalog⁴³ contains a table of universal and PEP security capabilities that apply across use cases, but not all apply to every use case. Each will contain a set of relevant security capabilities, based on agency pilot implementations and best practices. Additional security capabilities may be employed by agencies to reflect agency requirements, risk tolerances, and other factors. The EaaS guidance in the Cloud Use Case is one use case where some PEP security capabilities are not applicable.

⁴³ “Trusted Internet Connections webpage TIC Guidance Quick Links under Security Capabilities Catalog,” Cybersecurity and Infrastructure Security Agency. <https://www.cisa.gov/tic>.

For traceability, the security capabilities not included in this section of the use case are listed below by PEP capability group.

- Web: All
- Networking: All
- Resiliency: All
- Intrusion Detection: All
- DNS: Domain Name Sinkholing
- DNS: Domain Name Verification for Agency Clients
- DNS: CISA's Protective DNS Service
- Enterprise: Virtual Private Network
- Enterprise: Application Container
- Enterprise: Remote Desktop Access
- Services: All⁴⁴
- Unified Communication and Collaboration: All

Due to the unique security considerations for this use case, new security capabilities are included in the Email PEP group. These capabilities have been added to the current version of the Security Capabilities Catalog. The new security capabilities are detailed in the subsequent tables and listed here by PEP capability group for traceability.

- Email: Sender Denylisting
- Email: Post-Delivery Protections
- Email: Malicious File Protections
- Email: Adaptive Email Protections
- Email: Email Labeling
- Email: User Tipping
- Email: Content Filtering
- Email: User Digital Signatures for Outgoing Email
- Email: Encryption for Outgoing Email
- Email: Mail Content Query
- Email: Domain Reputation Protection

5.4.1 Universal Security Capabilities

Universal security capabilities are enterprise-level capabilities that outline guiding principles for TIC use cases and apply across all use cases. Agencies have the discretion to determine the level of rigor for applying universal security capabilities in accordance with federal guidelines and their risk tolerance.

When considering the universal security capabilities, agencies must understand what is provided by EaaS provider, what is required of the agency, what needs to be supplemented with an additional product or service, and how to integrate capabilities with their enterprise solution to fulfill each capability.

Table 14 provides a list of the universal security capabilities that apply to the EaaS guidance in the Cloud Use Case and implementation guidance for agencies to consider. Most agencies will have an existing enterprise solution for the universal security capabilities; as agencies deploy the EaaS guidance in the Cloud Use Case, the guidance below can be integrated into their existing solutions. While universal security capabilities are broadly applicable, the circumstances and threats associated with email in the cloud require agencies to consider the security challenges that may need to be addressed.

⁴⁴ Services PEP Security Capabilities are satisfied by the Email PEP Security Capabilities for this use case.

Table 14: Universal Security Capabilities for EaaS

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
Backup and Recovery	Backup and recovery entails keeping copies of configuration and data, as needed, to allow for the quick restoration of service in the event of malicious incidents, system failures, or corruption.	<i>Refer to Table 2.</i>
Central Log Management with Analysis	Central log management with analysis is the collection, storage, and analysis of telemetry, where the collection and storage are designed to facilitate data fusion and where the security analysis aids in discovery and response to malicious activity.	<i>Refer to Table 2.</i>
Configuration Management	Configuration management is the implementation of a formal plan, preferably automated, for documenting and managing changes to the environment and monitoring for deviations.	EaaS deployments often allow a degree of self-service by users to mailboxes, mailing lists, and other service resources. Agencies need to understand the self-service options that are available to their users, as well as any risks associated with their misconfiguration or inappropriate use. For example, allowing users to auto-forward email to external accounts may facilitate continued data exfiltration if the account is compromised, or facilitate shadow IT practices if users forward their work email to personal accounts.
Incident Response Plan and Incident Handling	Incident response planning and incident handling is the documentation and implementation of a set of instructions, procedures, or technical capabilities to sense and detect, respond to, limit consequences of malicious cyberattacks, and restore the integrity of the network and associated systems.	<p>Given the limited visibility into the EaaS environment, incident response will be a shared responsibility between the EaaS provider and the agency. The EaaS provider will be responsible for incident response handling for the infrastructure underlying the EaaS services, including services, operating systems, and all networks and hardware devices.</p> <p>Agencies should research the EaaS provider's incident response process, and make sure it aligns with their risk tolerances. In addition, agencies should be aware of how and when they will be notified in the event of an incident. The agency should guarantee that the EaaS provider notifies the agency within an acceptable amount of time when an incident impacts their data. Agencies should ensure that the EaaS deployment has mechanisms for globally searching all mailboxes and attachments to identify what mailboxes need to be sanitized in response to a data spillage.</p> <p>Agencies should include the agency email service in their incident response plans and review and update the plan if the EaaS provider updates its incident response policies. Agencies should plan for alternate ways to communicate during an incident or response in an event that involves their service or the EaaS provider. In other words, if an agency's email is compromised, there may need to be an alternate method for critical communications within the agency.</p>

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
Inventory	Inventory entails developing, documenting, and maintaining a current inventory of all systems, networks, and components so that only authorized devices are given access and unauthorized and unmanaged devices are found and restricted from gaining access.	Agencies should maintain awareness of all authorized users and mailboxes associated with the agency. Agencies should track how those user accounts and mailboxes are being accessed, including tracking the devices with access or synchronize user mailboxes, especially when user-furnished. Agencies should track which EaaS providers are authorized to receive email for agency domains and which services (e.g., EaaS providers, third-party mailers) are authorized to send email on behalf of the agency. Cryptographic certificates and keys (e.g., keys for DKIM) should also be carefully inventoried, where certain private keys may either be maintained by or need to be shared with the EaaS provider. Inventory of cryptographic keys should keep track of not only the keys, but who has had access to private keys.
Least Privilege	Least privilege is a design principle whereby each entity is granted the minimum system resources and authorizations that the entity needs to perform its function.	Agencies need to understand the authorizations available in the EaaS environment to ensure the accuracy of the least privilege in the EaaS deployment. EaaS deployments often require a degree of self-service by users to mailboxes, mailing lists, and other service resources. Agencies need to ensure that the self-service permissions for these resources align with least privilege needs.
Secure Administration	Secure administration entails securely performing administrative tasks by using secure protocols.	<i>Refer to Table 2.</i>
Strong Authentication	Strong authentication verifies the identity of users, devices, or other entities through rigorous means (e.g., MFA) before granting access.	<p>To enable compatibility with legacy email client applications, email providers often support protocols, like Post Office Protocol (POP) and Internet Message Access Protocol (IMAP), that have limited support for strong authentication mechanisms.⁴⁵ Agencies should understand the authentication techniques that can be used with legacy protocols and should disable protocols that do not support strong authentication in line with agency risk tolerances.</p> <p>Agencies allowing web-based access to their agency email service need to ensure that only strong authentication mechanisms can be used to access the service. The OMB Zero Trust Strategy (M-22-09) indicates that agencies should avoid “authentication methods that fail to resist phishing,” particularly methods supplying codes through Short Message Service, phone calls, or push notifications.⁴⁶</p> <p>Email providers allow users and client applications to request tokens that enable access to the service without requiring re-authentication. Agencies will need to understand the lifetimes of these tokens and the methods used to revoke them to ensure they can properly manage and revoke access to the service.</p>

⁴⁵ CISA Phishing Resistant Multifactor Authentication, www.cisa.gov/mfa.

⁴⁶ “Moving the U.S. Government Toward Zero Trust Cybersecurity Principles,” Office of Management and Budget M-22-09 (2022). <https://www.whitehouse.gov/wp-content/uploads/2022/01/M-22-09.pdf>.

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
Time Synchronization	Time synchronization is the coordination of system (e.g., servers, workstations, network devices) clocks to minimize the difference between system clock times and enable accurate comparison of timestamps between systems.	<i>Refer to Table 2.</i>
Vulnerability Management	Vulnerability management is the practice of proactively working to discover vulnerabilities by including the use of both active and passive means of discovery and by taking action to mitigate discovered vulnerabilities.	<i>Refer to Table 2.</i>
Patch Management	Patch management is the identification, acquisition, installation, and verification of patches for products and systems.	Agencies should consider patch management holistically, including the email service, security capabilities, and the email clients. Agencies should consider requiring device health checks and security posture before granting access to the EaaS provider to ensure client versions are in alignment with security policies.
Auditing and Accounting	Auditing and accounting include capturing business records (e.g., logs and other telemetry), making them available for auditing and accounting as required, and designing an auditing system that considers insider threat (e.g., separation of duties violation tracking) such that insider abuse or misuse can be detected.	<i>Refer to Table 2.</i>
Resilience	Resilience entails ensuring that systems, services, and protections maintain acceptable performance under adverse conditions.	<i>Refer to Table 2.</i>
Enterprise Threat Intelligence	Enterprise threat intelligence is the usage of threat intelligence from private or government sources to	<i>Refer to Table 2.</i>

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
	implement mitigations for the identified risks.	
Situational Awareness	Situational awareness is maintaining effective current and historical awareness across all components.	<p>Agencies need to ensure that their EaaS deployments are integrated into their overall situational awareness tools and processes. Agencies should consider, where possible, integrating telemetry from their EaaS deployments into centralized situational awareness tools to ensure a holistic view across the enterprise.</p> <p>Automated reporting can be a key part of maintaining situational awareness, and email is a common reporting mechanism. Agencies need to consider the effect that a compromise or inaccessibility of the agency email service might have on their situational awareness and the potential mitigations that might allow for continuity of reporting and situational awareness.</p>
Dynamic Threat Discovery	Dynamic threat discovery is the practice of using dynamic approaches (e.g., heuristics, baselining) to discover new malicious activity.	Agencies should ensure that their dynamic threat discovery solutions enable detection of anomalous user activity in the email service, including search or email downloads, email sending or forwarding, and changes to email forwarding rules or policies.
Policy Enforcement Parity	Policy enforcement parity entails consistently applying security protections and other policies, independent of the communication mechanism, forwarding path, or endpoints used.	To ensure that users have a consistent set of policies independent of their client or access location, agencies should consider policy enforcement positioning either at the email service or in front of the email service using a service like a SECaaS. Additionally, agencies should ensure that a commensurate set of policies are applied independent of the mechanisms that users employ to access the email service (e.g., Messaging Application Programming Interface [MAPI], web, and web API).
Effective Use of Shared Services	Effective use of shared services means that shared services are employed, where applicable, and individually tailored and measured to independently validate service conformance and offer effective protections for tenants against malicious actors, both external and internal to the service provider.	<i>Refer to Table 2.</i>
Integrated Desktop, Mobile, and Remote Policies	Integrated desktop, mobile, and remote policies define and enforce policies that apply to a given agency entity independent of its location.	<i>Refer to Table 2.</i>

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
User Awareness and Training	User awareness and training entails that all users are informed of their roles and responsibilities and that appropriate cybersecurity education is provisioned to enable users to perform their duties in a secure manner.	<p>All users should be trained in email security best practices and email security awareness in order to reduce the number of email data leaks and to prevent email threats such as phishing, malware, and malicious links.</p> <p>Users should be trained to recognize and avoid suspicious emails (e.g., phishing, and various types of social engineering) to reduce agency users' susceptibility to phishing and spear phishing. Educated users can sometimes detect and avoid malicious spam that is not detected automatically.</p> <p>Users should be trained in what agency data can and cannot be sent in email and how to protect any agency data that is sent in email, commensurate with agency policy and risk. Agencies should continue to update phishing awareness training and email security training as new threats emerge and email attacks become more sophisticated.</p> <p>Additionally, administrators and related IT staff should have the necessary training to manage, support, and secure agency email services.</p>
Supply Chain Risk Management	Supply chain risk management involves implementing a systematic process for managing cyber supply chain risk exposures, threats, and vulnerabilities throughout the supply chain and developing risk response strategies to the risks presented by the supplier, the supplied products and services, or the supply chain.	Agencies should consider and implement supply chain risk management prior to the acquisition of an EaaS solution. When agencies add third-party products or services into their EaaS solution, they should consider the supply chain of any third-party products or services.
Resource Lifecycle Management	Resource lifecycle management is the end-to-end process of managing resources from development to operation to retirement, such that resources are provisioned and decommissioned in conjunction with the applications they support.	As agencies transition email services into cloud-based deployments, they should take some unique factors into consideration, in addition to the other resource lifecycle management issues that come with cloud. Email accounts progress through a natural cycle between creation and deletion, and agencies should ensure they have processes in place for both creation and deactivation of email accounts. Cloud email storage is unique because the files will be stored remotely, and the administrators may not control the underlying hardware. Agencies should ensure they are able to enforce policies for auto-archival of email, inbox and individual email size limits, and time-based deletion. Upon account termination, agencies should ensure policies are in place to retain required email files and logs, to prohibit receipt to deactivated emails, and ensure email account activation status is in sync with the status of a tied user account. Agencies must comply with the Federal Records Act, observing the appropriate retention periods for messages in accordance with National Archive and Record Administration's Bulletin 2013-02 on 'Capstone Approach to Managing Email' or successor requirements.
Security Test and Exercise	Security tests (e.g., penetration testing or red teaming) verify the	Given the role of email in security monitoring and incident response, agencies need to account for the email service in their security testing and exercise procedures and how a compromise of the email service can

Capability	Description	Use Case Guidance and Deployment-Specific Guidance
	extent to which a system resists active attempts to compromise its security. Security exercises are simulations of emergencies that validate and identify gaps in plans and procedures.	be used to limit visibility into threat actor activities. Agency exercise procedures should include exercises that include scenarios where access to the agency email service is limited to ensure response capability when the email service has been compromised or made unavailable.
Continuous Monitoring Reporting	Continuous monitoring reporting entails the maintenance of ongoing awareness of informational security, vulnerabilities, and threats to support organizational risk management decisions.	<i>Refer to Table 2.</i>
Governance and Policy Auditing	Governance and policy auditing entails validating the proper definition, application and enforcement of agency rules and policies.	<i>Refer to Table 2.</i>

5.4.2 Policy Enforcement Point Security Capabilities

PEP security capabilities are primarily focused on the network level and inform technical implementation for a given use case, such as communication with agency-sanctioned CSPs. Agencies have the discretion to determine the applicability and level of rigor necessary for applying PEP security capabilities based on the specific cloud service deployed, the policy enforcement options available, federal guidelines, and risk tolerance. From the Security Capabilities Catalog, the PEP security capability groups applicable to the EaaS guidance in the Cloud Use Case correspond to the following security functions:

- Files
- Email
- DNS
- Enterprise
- Data Protection
- Identity

Agencies may determine the rigor of the security capabilities commensurate with risk and in accordance with federal guidelines, while considering mission needs and available policy enforcement options.

Security capabilities that are not applicable to this use case are listed at the beginning of Section 5.4. The PEP security capability listing is not exhaustive. Additional security capabilities may be deployed by agencies to reflect their risk tolerances, early adoption of security capabilities, the maturity level of existing cyber programs, etc.

5.4.2.1 Files PEP Security Capabilities

Agencies need to ensure that the applied file protection capabilities align with agency threats. Agencies should apply file protection capabilities to all incoming email, including attachments and body content. Agencies should, where available, consider protections that allow for determinations to be made on receipt as well as retroactively or during attempted access to ensure protection against files found to be malicious after receipt. When files are detected, the full email or the individual file attachments may be quarantined. File protections may misidentify legitimate files and agency users should have methods for accessing misidentified files. Table 15 lists the applicable Files PEP Security Capabilities for the EaaS guidance in the Cloud Use Case.

Table 15: Files PEP Security Capabilities for EaaS

Capability	Description	Use Case Guidance
Anti-malware	Anti-malware protections detect the presence of malicious code and facilitate its quarantine or removal.	Agencies should apply anti-malware protections to all incoming email, including attachments and body content. Agencies should consider applying anti-malware protections to outgoing email to detect the potential use of compromised accounts.
Content Disarm and Reconstruction	Content disarm and reconstruction technology detects the presence of unapproved active content and facilitates its removal.	<p>Agencies may consider employing content disarm and reconstruction technologies to incoming email, including attachments and body content, to decrease the attack surface across all agency user device types.</p> <p>Content disarm and reconstruction technologies may change documents in ways that render them unsuitable for agency use. Agencies should consider options for making the original file available to agency users on an as-needed basis. Agencies may also employ methods for agency users to access unmodified files from trusted sources.</p>
Detonation Chamber	Detonation chambers facilitate the detection of malicious code using protected and isolated execution environments to analyze the files.	<i>Refer to Table 3.</i>
Data Loss Prevention	DLP technologies detect instances of the exfiltration, either malicious or accidental, of agency data.	Email DLP solutions can protect against the malicious or accidental exfiltration of sensitive agency data. Agencies should consider applying email DLP solutions for all outgoing email. Agencies may consider applying similar capabilities to email received from external entities to ensure that data received by and stored on the agency email service aligns with agency data requirements and risk tolerance. Agencies need to ensure that DLP solutions can detect agency data. As entities can easily send encrypted email that limits the visibility of DLP solutions, agencies need to understand the protections available from DLP solutions for encrypted traffic.

5.4.2.2 Email PEP Security Capabilities

Agencies may have multiple campus locations and remote users across the country and possibly abroad. Since of the distributed nature of agency users and partners, an agency's email service is an important communication tool for business operations. An agency's email service is a frequent target for adversaries and therefore, agencies should carefully consider how security capabilities are deployed for their agency email service. Many of the email security capabilities identified in this group are offered by EaaS providers. However,

agencies should consider each capability with respect to security controls offered by the EaaS provider and use their risk profile to align the security controls, potentially augmenting them with additional capabilities. Table 16 lists the applicable Email PEP Security Capabilities for the EaaS guidance in the Cloud Use Case.

Table 16: Email PEP Security Capabilities for EaaS

Capability	Description	Use Case Guidance
Anti-phishing Protections	Anti-phishing protections detect instances of phishing and prevent users from accessing them.	<p>Various phishing techniques may be used against agency users and can often be tailored to the specific types of agency communications. Agencies should understand the threats applicable to them and tailor anti-phishing protections to those threats.⁴⁷ Agencies should consider anti-phishing protections that integrate ML techniques to understand the types of emails that users receive and tailor the phishing protections accordingly.</p> <p>Anti-phishing capabilities often employ the results of domain authentication techniques (e.g., DKIM, SPF, and Domain-based Message Authentication Reporting and Conformance [DMARC]) and anti-spam protections, and agencies should understand how the anti-phishing capabilities integrate this information, especially when these capabilities are provided by separate vendors.</p> <p>Agencies should consider the risk of spear phishing when assigning agency email addresses. Email addresses for particularly significant users or entity accounts (e.g., IT, human resources, chief information officers) should be quite distinct from any other email addresses, potentially following an obviously different pattern(s).</p>
Anti-spam Protections	Anti-spam protections detect and quarantine instances of spam.	Agencies should align anti-spam protections with the types of spam and business email that they receive. Anti-spam protections may misidentify legitimate business email as spam, and agency users should have methods for accessing misidentified emails. Agencies should know the options for designating specific senders or email as not being spam.
Authenticated Received Chain	Authenticated received chain allows for an intermediary, like a mailing list or forwarding service, to sign its own authentication of the original email, allowing downstream entities to accept the intermediary's authentication even if the email was changed.	Agencies may consider, where available, services that use authenticated receive chain to improve the accuracy of DMARC determinations in situations where agency users are receiving email that has been sent through a forwarding service (e.g., a mailing list).

⁴⁷ "Counter Phishing Recommendations for Federal Agencies," Cybersecurity and Infrastructure Security Agency (2020). https://www.cisa.gov/sites/default/files/publications/Capacity_Enhancement_Guide-Counter-Phishing_Recommendations_for_Federal_Agencies_1_0.pdf.

Capability	Description	Use Case Guidance
Data Loss Prevention	DLP technologies detect instances of the exfiltration, either malicious or accidental, of agency data.	Email DLP solutions can protect against the malicious or accidental exfiltration of sensitive agency data. Agencies should consider applying email DLP solutions for all outgoing email. Agencies may consider applying similar capabilities to email received from external entities to ensure that data received by and stored on the agency email service aligns with agency data requirements and risk tolerance. Agencies need to ensure that DLP solutions can detect agency data. As entities can easily send encrypted email that limits the visibility of DLP solutions, agencies need to understand the protections available from DLP solutions for encrypted traffic.
Domain Signature Verification for Incoming Email	Domain signature verification protections authenticate incoming email according to the DMARC email authentication protocol defined in Request for Comments (RFC) 7489. ⁴⁸	<p>Agencies need to understand how email that fails verification or cannot be verified (e.g., quarantine, rejection) and should, where feasible, consider methods to allow agency users to retrieve legitimate business email that could not be verified or failed verification. Agencies may need to work with the EaaS provider and external entities if legitimate business email consistently fails verification.</p> <p>Domain signature verification uses DNS resolution to retrieve external domain information as part of verification. Agencies should understand the DNS protections applied when the verification protections retrieve external domain information, and whether they can provide commensurate protections and visibility. Agencies need to ensure that their EaaS provider is enabling verification for incoming agency emails.⁴⁹</p>
Domain Signatures for Outgoing Email	Domain signature protections facilitate the authentication of outgoing email by signing the emails and ensuring that external parties may validate the email signatures according to the DMARC email authentication protocol that is defined in RFC 7489.	Agencies will need to understand the capabilities offered by the EaaS provider for signing outgoing emails, and should, when available, enable DKIM and DMARC. ⁵⁰ Agencies need to ensure that the signing certificates are managed using appropriate key lifecycle management ^{51 52} and understand how to revoke and update the certificates in the case of compromise. Agencies will need to coordinate the domain signature configurations between their DNS and email services.
Encryption for Email Transmission	Email services are configured to use encrypted connections, when possible, for communications between clients and other email servers.	Agencies should ensure that the EaaS provider uses encryption for email transmission, following security recommendations 5-2, 5-3, 5-4, and 7-1 of NIST SP 800-177, Revision 1. ⁵³ Agencies should ensure that their email service only allows agency clients to communicate with it using encrypted channels. ⁵⁴

⁴⁸ "Domain-based Message Authentication, Reporting, and Conformance Request for Comments: 7489," Internet Engineering Task Force (2015), <https://tools.ietf.org/html/rfc7489>.

⁴⁹ "Binding Operational Directive 18-01 Enhance Email and Web Security," Cybersecurity and Infrastructure Security Agency (2017). <https://www.cisa.gov/binding-operational-directive-18-01>.

⁵⁰ "Binding Operational Directive 18-01 Enhance Email and Web Security," Cybersecurity and Infrastructure Security Agency (2017). <https://www.cisa.gov/binding-operational-directive-18-01>.

⁵¹ "SP 800-177 Rev. 1, Trustworthy Email," National Institute of Standards and Technology (2019). <https://csrc.nist.gov/publications/detail/sp/800-177/rev-1/final>.

⁵² "SP 800-57 Part 1 Rev. 5, Recommendation for Key Management: Part 1 – General," National Institute of Standards and Technology (2020). <https://csrc.nist.gov/publications/detail/sp/800-57-part-1/rev-5/final>.

⁵³ "SP 800-177 Rev. 1, Trustworthy Email," National Institute of Standards and Technology (2019). <https://csrc.nist.gov/publications/detail/sp/800-177/rev-1/final>.

⁵⁴ "Binding Operational Directive 18-01 Enhance Email and Web Security," Cybersecurity and Infrastructure Security Agency (2017). <https://www.cisa.gov/binding-operational-directive-18-01>.

Capability	Description	Use Case Guidance
Malicious Link Protections	Malicious link protections detect malicious links in emails and prevent users from accessing them.	Agencies should, where available, consider malicious link protections that allow for malicious link determinations to be made on receipt as well as retroactively or during attempted access to provide protection against links that are found to be malicious after receipt.
Link Click-through Protections	Link click-through protections ensure that when a link from an email is clicked, the requester is directed to a protection that verifies the security of the link destination before permitting access.	Agencies need to ensure that link click-through protections do not interfere with one-time use links (e.g., password reset). Agencies may consider combinations of link click-through protections (e.g., static checks coupled with browser isolation).
EINSTEIN 3 Accelerate (E3A) Email Protections	EINSTEIN 3 Accelerate (E3A) is an intrusion prevention capability, provided by NCPS, provided by CISA, that includes an email filtering security service.	Agencies may need to work with CISA to ensure commensurate protections and visibility are available when EaaS deployments do not integrate E3A protections.
Sender Denylisting	Sender denylisting protections prevent the reception of email from denylisted senders, domains, or email servers.	Agencies should consider using enterprise threat intelligence, potentially cloud-native, that help automate denylisting for senders, sending domains or addresses that are known or suspected malicious. Agencies should consider sender denylist solutions that can be applied retroactively to received email for senders subsequently included in the denylists.
Post-Delivery Protections	Post-delivery protections apply updated email protections to already delivered emails, enabling quarantining and mitigation for emails in mailboxes.	Agencies should consider protections for email that can be applied both on receipt and retroactively to received email. These could be deployed as a single protection that applies to both, or as distinct protections. If separate protections are employed, agencies need to understand any differences in the protections (see Mail Content Query).
Malicious File Protections	Malicious file protections detect malicious attachment files in emails and prevent users from accessing them.	Agencies should consider malicious file protections that apply both static and dynamic analysis techniques. Malicious file protections should allow for malicious files to be detected both on receipt as well as retroactively when updated determinations can be made.
Adaptive Email Protections	Adaptive email protections involve employing risk-based analysis in the application and enforcement of email protections.	As spear phishing techniques become more sophisticated, agencies should consider using a risk-based approach to the application of email protections. This may include the use of user profile and group profile awareness when applying and enforcing email protections. For example, senior leadership at agencies may be a continuous target of threat actors, and as a result, an agency may consider the protections in place for these users.
Email Labeling	Email labeling is the process of automatically tagging incoming or outgoing email to manage risk.	Agencies should deploy automated email tags and banners to email subject lines and/or bodies so that the tag can inform users of potential risks about the email. These tags, labels, and banners may be generated by other email security capabilities (e.g., content filtering). Examples of email labeling include distinguishing between internal and external email or labeling potential spam.

Capability	Description	Use Case Guidance
User Tipping	User tipping capabilities enable users to report emails, attachments, or links they suspect to be phishing attempts, spam, or otherwise malicious.	User tipping, also called “email reporting,” enables agencies to allow users to report potentially malicious emails, potentially through an EaaS-native capability, a dedicated email address, a webpage, or other means. Where possible, the user tipping should provide enough context, including user supplied context, to enable an accurate understanding of the tip. Agencies should also understand opportunities for tipping information to the EaaS provider as well as to any services providing security capabilities for their agency email service.
Content Filtering	Content filtering protections detect the presence of unapproved content and facilitate its removal or denial of access.	Agencies content filtering protections for email comprise a variety of policies, including blocking unauthorized or illegal content, removing common email tracking mechanisms, and allowing or permitting certain attachment file types. Agencies should ensure that available content filter protections align with their policy needs and requirements. Agencies should consider mechanisms that enable agency users to access the filtered content to account for potential accommodations for legitimate uses that are impacted by the filtering policies.
User Digital Signatures for Outgoing Email	User digital signature protections enable users to digitally sign their emails, allowing external parties to authenticate the email's sender and its contents.	Agencies should consider EaaS providers that allow for users to sign their emails (e.g., using S/MIME). To ensure alignment with agency needs and risk tolerance, agencies will need to understand the methods available for signing emails and the impacts they might have on other security protections. Agencies need to ensure that user signing keys are managed using appropriate key lifecycle management and understand how to revoke and update the keys in the case of compromise. ^{55 56}

⁵⁵ “SP 800-177 Rev. 1, Trustworthy Email,” National Institute of Standards and Technology (2019).

<https://csrc.nist.gov/publications/detail/sp/800-177/rev-1/final>.

⁵⁶ “SP 800-57 Part 1 Rev. 5, Recommendation for Key Management: Part 1 – General,” National Institute of Standards and Technology (2020). <https://csrc.nist.gov/publications/detail/sp/800-57-part-1/rev-5/final>.

Capability	Description	Use Case Guidance
Encryption for Outgoing Email	Email encryption protections allow for the encryption of outgoing emails, limiting the visibility of their contents to the intended recipients.	<p>Agencies should ensure that the EaaS provider uses encrypted communications when sending outgoing email to external email services. Agencies need to ensure that user encryption keys are managed using appropriate key lifecycle management and understand how to revoke and update the keys in the case of compromise.</p> <p>Agencies should consider EaaS providers that allow for end-to-end encryption of email content and attachments. End-to-end encryption mechanisms often enable users to encrypt outgoing email, and agencies need to understand the impact it may have on visibility for outgoing email protections. Some EaaS providers may allow the option of storing the keys needed to encrypt the outgoing email. This functionality can ease users' use of encryption, especially when they use multiple devices to send email, and can enable security capability visibility into emails. However, agencies will need to consider the security implications of providing these keys to the EaaS provider as the keys can be used to authenticate messages more strongly as having come from the agency users.</p> <p>Some EaaS providers can provide functionality that enables encrypted content to be sent to external entities that may not support receiving encrypted emails, often sending an unencrypted email to the external entity with a link to the encrypted content. As these links can be used to access the encrypted data, agencies need to ensure that emails are transmitted only over encrypted channels and that the access to this encrypted content is provided only over protected channels and for limited timeframes.</p>
Mail Content Query	Mail content query enables search and discovery for email across agency mailboxes.	Agencies should consider integrating email search and discovery into their incident response procedures to enable determination of all instances of a malicious email, inappropriate email, or data breach.
Email Domain Reputation Protection	Email domain reputation protection entails monitoring an email domain's reputation and employing policies to help protect the email domain's reputation.	Agencies should enable email authentication solutions (e.g., DMARC, SPF, DKIM) to help ensure that external email services can authenticate genuine agency email. Agencies should understand the email sent on their behalf and ensure that it aligns with agency policy and with best practices for maintaining effective domain reputation.

5.4.2.3 Domain Name System PEP Security Capabilities

DNS provides a key underpinning of agency email communication, enabling both the sending and receiving of email with external entities. When sending email to an external email service, the agency email service will use DNS to look up the email service associated with the given domain along with information. When an external entity sends an email to the agency, their email service will use DNS to look up and validate the agency email service. Additionally, when the agency email receives the email from the external entity, the email service will use DNS to look up information needed to verify the sender of the email and the validity of the received email contents. Table 17 lists the applicable DNS PEP Security Capabilities for the EaaS guidance in the Cloud Use Case.

Table 17: Domain Name System PEP Security Capabilities for EaaS

Capability	Description	Use Case Guidance
Domain Name Validation for Agency Domains	Domain name validation protections ensure that all agency domain names are secured using DNSSEC, enabling external entities to validate their resolution to the domain names.	Agencies will need to store a variety of components in DNS services that each play a role in enabling them to communicate with external entities. For external entities to be able to send email to the agency, the agency needs to configure DNS records, called mail exchanger (MX) records, that provide external entities with the address of the agency email service. For the agency to be able to send email to external entities, they need to configure SPF records to define which email services are authorized to send email on behalf of the agency, DKIM records to allow the external entity to authenticate and validate emails received from agency email services, and DMARC records to ensure that external providers understand how to handle email received from unauthenticated or unauthorized sources. Agencies need to ensure that these components are available in DNS providers that support DNSSEC to ensure that external entities can validate the information they receive.
Domain Name Monitoring	Domain name monitoring allows agencies to discover the creation of or changes to agency domains.	Agencies should ensure domain name monitoring solutions integrate all the various domain information, including MX, SPF, DKIM, and DMARC, used by external entities to communicate with the agency email service.

5.4.2.4 Enterprise PEP Security Capabilities

Email forms a core component of agency environments, handling everything from internal and external communication to alerting and status monitoring. Agencies need to understand how to integrate their agency email service into their overall workflows, accounting for the effect of loss of connectivity to ensure continuity of operations when access to the agency email service is interrupted. Table 18 lists the applicable Enterprise PEP Security Capabilities for the EaaS guidance in the Cloud Use Case.

Table 18: Enterprise PEP Security Capabilities for EaaS

Capability	Description	Use Case Guidance
Security Orchestration, Automation, and Response	SOAR tools define, prioritize, and automate the response to security incidents.	<p>Agencies will need to take a holistic view to best understand how to integrate the agency email service into their overall SOAR infrastructure. They may be able to integrate their EaaS deployment directly into their existing SOAR solutions. Alternatively, they may consider a SOAR solution tailored for their agency email service so long as the detections and responses available in the solution align with overall agency SOAR needs. When employing a different solution, agencies need to understand the differences between their existing SOAR solution and the solution tailored for the agency email service to maintain an enterprise wide understanding of their security protections.</p> <p>Given the ease of quickly exfiltrating data from email environments, agencies should consider solutions that enable automatic responses to malicious activity, including user-centric responses like account disabling and email-centric responses like quarantining.</p>

Capability	Description	Use Case Guidance
Shadow Information Technology Detection	Shadow IT detection systems detect the presence of unauthorized software and systems in use by an agency.	<p>Agencies should consider methods to ensure that agency users' access to personal email accounts aligns with agency risk tolerance. Agencies may need to take a holistic approach as access to personal email accounts can occur over traditional email protocols, as well as through web-based email methods. As EaaS providers may also host personal email accounts, agencies may need to be able to differentiate personal use from business use.</p> <p>Agencies should ensure that the devices that agency users utilize to access the agency email service align with agency-sanctioned endpoint policies.</p>
Costs Monitoring	Costs monitoring entails the monitoring of costs incurred by enterprise resources.	<i>Refer to Table 9.</i>

5.4.2.5 Data Protection PEP Security Capabilities

Data protection is the process of maintaining the confidentiality, integrity, and availability of an agency's data consistent with their risk management strategy. It is important that agencies secure their data from corruption, compromise, or loss. Agencies should have processes and tools in place to protect agency data, prevent data exfiltration, and ensure the privacy and integrity of data, considering that data may be accessed from devices beyond the protections and perhaps administration of agencies. Agencies do not have control over physical protections for email stored at an EaaS provider. Therefore, the application of data protection security capabilities is critical to securing agency email in its cloud deployment. Agencies should consider the sensitivity of data when applying rigor to these Data Protection PEP Security Capabilities. Agencies should ensure that policies, procedures, and incident response are adapted to accommodate email storage and use.

Table 19 lists the applicable Data Protection PEP Security Capabilities for the EaaS guidance in the Cloud Use Case.

Table 19: Data Protection PEP Security Capabilities for EaaS

Capability	Description	Use Case Guidance
Access Control	Access control technologies allow an agency to define policies concerning the allowable activities of users and entities to data and resources.	<p>Agencies should ensure that email is only accessed by authorized users using MFA, and that least privilege is enforced. Additionally, agencies should ensure that email is only accessed from trusted devices.</p> <p>Agencies should have policies in place to determine who can read users' emails and under what circumstances. Agencies should use data access controls that align with these policies. For example, email administrators should not be able to read any email processed by the email server; however, in the event of a cyber investigation, an agency may allow limited access to emails.</p>
Protections for Data at Rest	Data protection at rest aims to secure data stored on any device or storage medium.	Agencies should protect user mailboxes and file stores commensurate with their risk tolerance level. For maximum security, email should be stored encrypted. Cryptographic keys used for encrypting data in persistent storage (e.g., in mailboxes) should be different from keys used for the transmission of email messages.

Capability	Description	Use Case Guidance
Protections for Data in Transit	Data protection in transit, or data in motion, aims to secure data that is actively moving from one location to another, such as across the internet or through a private enterprise network.	In line with Executive Order 14028, ⁵⁷ agencies should protect email in transit so that the email is not modified in transit or sensitive data is not leaked. Agencies can use encrypt email transfer between servers or use end-to-end email encryption.
Data Loss Prevention	DLP technologies detect instances of the exfiltration, either malicious or accidental, of agency data.	Email DLP solutions can protect against the malicious or accidental exfiltration of sensitive agency data. Agencies should consider applying email DLP solutions for all outgoing email. Agencies may consider applying similar capabilities to email received from external entities to ensure that data received by and stored on the agency email service aligns with agency data requirements and risk tolerance. Agencies need to ensure that DLP solutions are able to detect agency data. As entities can easily send encrypted email that limits the visibility of DLP solutions, agencies need to understand the protections available from DLP solutions for encrypted traffic.
Data Access and Use Telemetry	Data access and use telemetry identifies agency-sensitive data stored, processed, or transmitted, including those located at a service provider, and it enforces detailed logging for access or changes to sensitive data.	An agency should track all access to agency email accounts and mailboxes. Most EaaS providers have native capabilities for logging, monitoring, and analyzing email access telemetry.
Data Labeling	The process of tagging data by categories to protect and control the use of data and identifying a level of risk associated with the data.	Email is the most common communication platform for all internal and external communications, and emails are often stored in agency's mailboxes for very long periods of time; therefore, an agency should make sure that email labeling (see Email: Email Labeling) is integrated into the agency's data labeling procedures. Agencies may consider the labeling of attachments in addition to messages.
Data Inventory	Inventory entails developing, documenting, and maintaining a current inventory of agency data.	Agencies should have an inventory of all user mailboxes and have procedures for removing mailboxes when a user no longer needs access to the agency's email service.

5.4.2.6 Identity PEP Security Capabilities

Strong verification of identity is a key component to EaaS, as agency users often access these services from remote locations with more limited visibility into user devices and environments coupled with the high potential for account compromise and data exfiltration. Agencies need to employ protections beyond simple identity authentication including checking for device security posture and compliance, and detecting anomalous or suspicious user behavior.

Email is a primary means of communication between agencies and external entities. Agencies need to ensure the identity of the email service is properly configured and securely managed to enable external entities to

⁵⁷ "Executive Order 14028 Improving the Nation's Cybersecurity," Office of Management and Budget (May 2021). <https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/>.

validate the email they receive as having come from the agency email service. Additionally, the email service identity enables agency entities to ensure that they are accessing the agency email service. Table 20 lists the applicable Identity PEP Security Capabilities for the EaaS guidance in the Cloud Use Case.

This capability group and all capabilities in Table 20 are new and will be added to the next version of the Security Capabilities Catalog during the next revision cycle.

Table 20: Identity PEP Security Capabilities for EaaS

Capability	Description	Use Case Guidance
Adaptive Authentication	Adaptive authentication aligns the strength of the user or entity authentication mechanisms to the level of risk associated with the requested authorization.	Email services enable users and client applications to request tokens that enable access without requiring re-authentication. Agencies will need to understand the accesses permitted by these tokens to ensure they cannot be used to bypass adaptive authentication controls.

Capability	Description	Use Case Guidance
Service Identity	Service identity ensures that users and entities can authenticate the identities of agency services.	<p>Identities for email services consist of a few different components, each of which plays a role in enabling the authentication of the service at various stages in the email sending and delivery workflows. When agencies send email to external email services, the external services may use the SPF to determine which addresses are authorized to send email from the agency. Agencies need to ensure that the appropriate SPF records are configured to enable this external validation and will need to maintain these records as changes occur in the agency environment to ensure that only addresses for agency-authorized services are included in the records.</p> <p>External email services use DKIM authentication to authenticate and validate emails received from agency email services. Under DKIM, the external email service retrieves a key for the email service using DNS. The agency email service then uses that key to sign outgoing email, enabling the external email service to verify the authenticity of the agency email service. Agencies will need to ensure that appropriate key lifecycle management^{58 59} for their domain signing keys, and the alignment of their DKIM configuration across DNS and the email services. Additionally, agencies should enable DMARC configuration to ensure that external providers understand how to handle email received from unauthenticated or unauthorized sources.</p> <p>Email services will also have TLS certificates to enable encryption and authentication of traffic between the email service and agency clients as well as being the email service and external email services. These may include certificates for the email protocols (e.g., MAPI, Extended Simple Mail Transport Protocol (ESMTP/TLS)), as well as for any web-based email services. Agencies need to ensure that these certificates are managed using appropriate key lifecycle management and understand how to revoke and update the certificates in the case of compromise.</p> <p>External entities use DNS to look up where to send email destined for agency domains. These DNS records, called MX records, provide external entities with the address of the email services where they should send the emails. Agencies need to ensure that these MX records are properly configured and are made available using DNS services that provide DNSSEC to enable the external entities to validate the records.</p> <p>If an agency needs to allow external entities to send messages on their behalf (e.g., mailing lists, newsletters), agencies need to understand how best to enable that functionality in alignment with the above service identity considerations and ensure that the change in security posture aligns with their risk tolerance.</p>
Entitlement Inventory	Entitlement inventory entails developing, documenting, and maintaining a current inventory of user and entity permissions and authorizations to agency resources.	Agencies need to understand the authorizations available in the EaaS deployment to ensure an accurate inventory. Agencies should consider methods of integrating these EaaS entitlements into their enterprise entitlement inventory to ensure a holistic understanding of entitlements. Agencies need to ensure that entitlement inventories track changes where agency users or entities can make changes to entitlements.

⁵⁸ "SP 800-177 Rev. 1, Trustworthy Email," National Institute of Standards and Technology (2019).

<https://csrc.nist.gov/publications/detail/sp/800-177/rev-1/final>.

⁵⁹ "SP 800-57 Part 1 Rev. 5, Recommendation for Key Management: Part 1 – General," National Institute of Standards and Technology (2020). <https://csrc.nist.gov/publications/detail/sp/800-57-part-1/rev-5/final>.

Capability	Description	Use Case Guidance
Secrets Management	Secrets management entails developing and using a formal process to securely track and manage digital authentication credentials, including certificates, passwords, and API keys.	<p>Email services contain a variety of keys that are used by external parties to verify the identity of the service and the validity of emails received from the agency. Additionally, agencies may support agency users and entities having keys that enable them to sign outgoing email and to send or receive encrypted email. Agencies need to ensure that these keys are managed using appropriate key lifecycle management processes and understand how to revoke them in cases of compromise. Agencies need to track where these keys are deployed to better understand the opportunities for key compromise.</p> <p>Email services enable users and client applications to request tokens that enable access to the service. Agencies will need to understand and track the creation and lifecycle of these tokens.</p>
Behavioral Baselineing	Behavioral baselining is capturing information about user and entity behavior to enable dynamic threat discovery and facilitate vulnerability management.	<p>Email deployments often allow for a variety of methods to access the service, including full-service applications on computers, mobile applications, and web-based access. Agency users may employ a combination of these methods, and behavioral baselining methods need to understand how the users access the service, and potentially how their behavior differs across access methods.</p> <p>Behavioral baselining methods need to account for the variety of activities that users can perform for the email service, including behaviors related to data access that might be evidence of data exfiltration, and behaviors related to sending emails which might be evidence of internal spear phishing. Additionally, behavioral baselining may need to account for changes in automated activities, like forwarding rules or other automated responses that are not directly invoked by user client activity.</p>
Enterprise Identity, Credential, and Access Management	Enterprise ICAM entails maintaining visibility into agency identities across agency environments and managing changes to those identities through a formal (preferably automated) process.	<p>Agencies that integrate ICAM across their enterprise environment need to understand the potential that opens for lateral movement into the agency email environment and the associated opportunities for access to agency email and exfiltration.</p> <p>Agency email identities may not map exactly to agency enterprise identities, with some agency enterprise identities not having associated email access or some identities that are specific to the agency email service. Agencies need to ensure they have a holistic understanding of identities, tracking the identities that exist, independent of where they originate.⁶⁰</p>
Multi-factor Authentication	MFA entails using two or more factors to verify user or entity identity.	<p>To enable compatibility with legacy email client applications, email providers often support protocols that do not support multifactor authentication. Agencies should, where feasible, disable protocols that do not support multifactor authentication, in line with agency risk tolerances.</p> <p>Additionally, email services enable users and client applications to request tokens that can be used to access the service without requiring MFA. Agencies will need to understand the access permitted by these tokens and ensure that the token lifetimes and ability to be revoked align with their risk tolerance.</p>

⁶⁰ Office of Management and Budget. "Enabling Mission Delivery through Improved Identity, Credential, and Access Management" (May 2019), <https://www.whitehouse.gov/wp-content/uploads/2019/05/M-19-17.pdf>.

Capability	Description	Use Case Guidance
Continuous Authentication	Continuous authentication entails validating and re-authenticating identity through the lifecycle of entity interactions.	Email services enable users and client applications to request tokens that enable access without requiring re-authentication. Agencies will need to understand the access permitted by these tokens and ensure that the token lifetimes align with their risk tolerance.

5.5 TELEMETRY REQUIREMENTS

As agencies transition from on-premises deployments to deployments in cloud environments, visibility by CISA must be preserved through information sharing. Figure 18 shows the conceptual architecture of the EaaS guidance in the Cloud Use Case, with the telemetry requirements added as lines on certain data flows. These lines, depicted in Figure 18, indicate when an agency must share telemetry with CISA. Subject to applicable law, CISA may require internal telemetry to be collected. The requirements for sharing telemetry data with CISA are only applicable to the data flows between the agency email service and external entities. Consult the NCPS program⁶¹ and CDM program⁶² for further details.

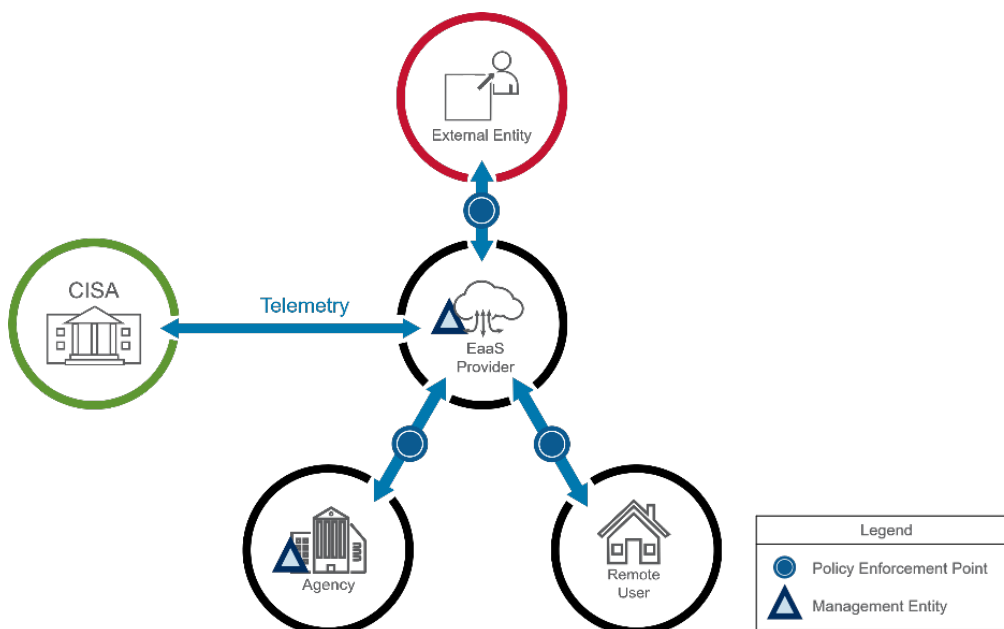


Figure 18: EaaS Telemetry Sharing with CISA

⁶¹ "National Cybersecurity Protection System," Cybersecurity and Infrastructure Security Agency. <https://cisa.gov/national-cybersecurity-protection-system-ncps>.

⁶² "Continuous Diagnostics and Mitigation," Cybersecurity and Infrastructure Security Agency. <https://cisa.gov/cdm>.

6. CONCLUSION

The *TIC 3.0 Cloud Use Case* defines how network and multi-boundary security should be applied in cloud environments. The use case is broken into two distinct parts, focusing on cloud deployments for:

1. IaaS, PaaS, and SaaS (Section 4)
2. EaaS (Section 5)

This document provides guidance on how an agency can configure its cloud data flows and apply relevant TIC security capabilities. Overall, the Cloud Use Case presents a total of eight network security patterns between guidance for (1) IaaS, PaaS, and SaaS, and (2) EaaS. This use case document should be used in conjunction with the Security Capabilities Catalog and any TIC overlays that are applicable to service providers that an agency employs.

APPENDIX A – GLOSSARY AND DEFINITIONS

This glossary contains terms and definitions that are used across the TIC documents and are not necessarily applicable to all use cases.

Boundary: A notional concept that describes the perimeter of a zone (e.g., mobile device services, general support system [GSS], Software-as-a-Service [SaaS], agency) within a network architecture. The bounded area must have an information technology (IT) utility.

Internet: The internet is discussed in two capacities throughout TIC documentation:

1. A means of data and IT traffic transport.
2. An environment used for web browsing purposes, referred to as “web.”

Managed Trusted Internet Protocol Services (MTIPS): Services under GSA’s Enterprise Infrastructure Solutions (EIS) contract vehicle that provide TIC solutions to government clients as a managed security service. It is of note that the EIS contract is replacing the GSA Network contract vehicle that is set to expire in Fiscal Year (FY) 2023.

Management Entity (MGMT): An entity that oversees and controls security capabilities. The entity can be an organization, network device, tool, service, or application. The entity can control the collection, processing, analysis, and display of information collected from the policy enforcement points (PEPs), and it allows IT professionals to control devices on the network.

National Cyber Protection System (NCPS): An integrated system-of-systems that delivers a range of capabilities, including intrusion detection, analytics, intrusion prevention, and information sharing capabilities that defend the civilian Federal Government’s information technology infrastructure from cyber threats. The NCPS capabilities, operationally known as EINSTEIN, are one of several tools and capabilities that assist in federal network defense.

Policy Enforcement Point (PEP): A security device, tool, function, or application that enforces security policies through technical capabilities.

Policy Enforcement Point Security Capabilities: Network-level capabilities that inform technical implementation for relevant use cases.

Reference Architecture: An authoritative source of information about a specific subject area that guides and constrains the instantiations of multiple architectures and solutions.

Risk Management: The program and supporting processes to manage information security risk to organizational operations (including mission, functions, image, and reputation), organizational assets, individuals, other organizations, and the United States. It includes: (1) establishing the context for risk-related activities; (2) assessing risk; (3) responding to risk once determined; and (4) monitoring risk over time.

Risk Tolerance: The level of risk or degree of uncertainty that is acceptable to organizations and is a key element of the organizational risk frame. An organization’s risk tolerance level is the amount of corporate data and systems that can be risked to an acceptable level.

Security Capability: A combination of mutually reinforcing security controls (i.e., safeguards and countermeasures) implemented by technical means (i.e., functionality in hardware, software, and firmware), physical means (i.e., physical devices and protective measures), and procedural means (i.e., procedures performed by individuals). Security capabilities help to define protections for information being processed, stored, or transmitted by information systems.

Security Pattern: Description of an end-to-end data flow between two trust zones. Security patterns may have an associated set of security capabilities or guidance to secure the data flow along with one or more of the zones.

Seeking Service Agency (SSA): An agency that obtains TIC services through an approved Multi-Service TICAP.

Security Information and Event Management (SIEM): An approach to security management that combines SIM (security information management) and SEM (security event management) functions into one security management system.

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

TIC: The term “TIC” is used throughout the Federal Government to denote different aspects of the TIC initiative; including the overall TIC program, a physical TIC access point (also known as a Traditional TIC), and a TICAP (see below). This document refers to TIC as an adjective or as the Trusted Internet Connections initiative.

TIC Access Point: The physical location where a federal civilian agency consolidates its external connections and has security controls in place to secure and monitor the connections.

TIC Access Provider (TICAP): An agency or vendor that manages and hosts one or more TIC access points. Single Service TICAPs serve as a TICAP only to their own agency. Multi-Service TICAPs also provide TIC services to other agencies through a shared-services model.

TIC Initiative: Program established to optimize and standardize the security of individual external network connections currently in use by the Federal Government, to include connections to the internet. Key stakeholders include CISA, OMB, and GSA.

TIC Overlay: A mapping from products and services to TIC security capabilities.

TIC Use Case: Guidance on the secure implementation and/or configuration of specific platforms, services, and environments. A TIC use case contains a conceptual architecture, one or more security pattern options, security capability implementation guidance, and CISA telemetry guidance for a common agency computing scenario.

Trust Zone: A discrete computing environment designated for information processing, storage, and/or transmission that dictates the level of security necessary to protect the traffic transiting in and out of a zone and/or the information within the zone.

Unified Communications and Collaboration: A collection of solutions designed to facilitate communication and collaboration, including in real-time, such as required by remote work or collaboration between locations.

Universal Security Capabilities: Enterprise-level capabilities that outline guiding principles for TIC use cases.

Web: An environment used for web browsing purposes. Also see Internet.

Zero Trust: A security model based on the principle of maintaining strict access controls and not trusting anyone by default, even those already inside the network perimeter.

APPENDIX B – RELATED FEDERAL GUIDELINES AND REQUIREMENTS

The following list of documents include the most recent version of the guidance documents available at the time of this publication, including drafts.

Cybersecurity and Infrastructure Security Agency, *Binding Operational Directive 18-01, Enhance Email and Web Security*, October 2017.

Cybersecurity and Infrastructure Security Agency, *Capacity Enhancement Guides for Federal Agencies: Implementing Strong Authentication*, October 2020.

Cybersecurity and Infrastructure Security Agency, *Cloud Security Technical Reference Architecture, Version 2.0*, June 2022.

Cybersecurity and Infrastructure Security Agency, *Cybersecurity Incident & Vulnerability Response Playbooks: Operational Procedures for Planning and Conducting Cybersecurity Incident and Vulnerability Response Activities in FCEB Information Systems*, November 2021.

Cybersecurity and Infrastructure Security Agency, *Zero Trust Maturity Model, Version 2.0*, April 2023.

Department of Defense, *Zero Trust Reference Architecture, Version 2.0*, July 2022.

Cybersecurity and Infrastructure Security Agency, *Extensible Visibility Reference Framework Guidebook Request for Comment Draft*, April 2022.

Federal Information Security Modernization Act of 2014 (P.L. 113-283), codified in relevant part in 44 U.S.C. §§ 3551-8.

National Institute of Standards and Technology Special Publication, 800-53, Revision 5, Security and Privacy Controls for Federal Information Systems and Organizations, December 2020.

National Institute of Standards and Technology Special Publication, 800-63-3, Digital Identity Guidelines, March 2020.

National Institute of Standards and Technology, Special Publication 800-177, Revision 1, Trustworthy Email, February 2019.

National Institute of Standards and Technology, Special Publication 800-207, Revision 1, Zero Trust Architecture, August 2020.

National Institute of Standards and Technology, Special Publication 800-210, General Access Control Guidance for Cloud Systems, July 2020.

Cybersecurity and Infrastructure Security Agency, *Secure Cloud Business Applications Microsoft 365 Baselines Draft*, December 2022.

Cybersecurity and Infrastructure Security Agency, *Secure Cloud Business Applications Technical Architecture Request for Comment Draft*, April 2022.

APPENDIX C – GLOSSARY FOR CLOUD USE CASE

This glossary contains cloud-specific terms and definitions that are used in this use case.

Cloud Access Security Broker (CASB): A Cloud-delivered Policy Enforcement Point that mediates access to cloud resources, enabling the application of security and other policies.

Cloud Infrastructure Entitlement Management (CIEM): Capabilities that facilitate the management of identities and entitlements in cloud and multi-cloud environments.

Cloud-Native Application Protection Platform (CNAPP): Capabilities that help align the visibility and security protections for deployed cloud applications.

Cloud Security Posture Management (CSPM): Capabilities that facilitate monitoring in cloud and multi-cloud environments by identifying, alerting on, and mitigating cloud vulnerabilities. Some CSPM capabilities that focus on managing and securing SaaS applications may be referred to as SaaS Security Posture Management (SSPM) solutions.

Cloud Workload Protection Platforms (CWPP): A platform designed to help facilitate visibility and management of security controls in cloud and multi-cloud environments, commonly including functions like system hardening, vulnerability management, host-based segmentation, system integrity monitoring, and application allow lists.

Desktop-as-a-Service (DaaS): A cloud computing offering where a CSP delivers cloud-hosted virtual desktops to end users in an organization. The CSP provides maintenance, back-up, updates, and data storage.

Email-as-a-Service (EaaS): A service provided to the consumers with tools to host email with unlimited storage and back up options.

Infrastructure-as-a-Service (IaaS): A service provided to the consumer for provision processing, storage, networks, and other fundamental computing resources where the consumer can 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).

Platform-as-a-Service (PaaS): A service provided to the consumer 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.

Security-as-a-Service (SECaaS): A Cloud-delivered service offering that provides cybersecurity capabilities and protections.

Software-as-a-Service (SaaS): A service provided to the consumer 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 (runs from resources stored on a central server instead of a localized hard drive) 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.