



EXECUTIVE OFFICE OF THE PRESIDENT
OFFICE OF MANAGEMENT AND BUDGET
WASHINGTON, D.C. 20503

January 26, 2022

M-22-09

MEMORANDUM FOR THE HEADS OF EXECUTIVE DEPARTMENTS AND AGENCIES

FROM: Shalanda D. Young
Acting Director

SUBJECT: Moving the U.S. Government Toward Zero Trust Cybersecurity Principles

This memorandum sets forth a Federal zero trust architecture (ZTA) strategy, requiring agencies to meet specific cybersecurity standards and objectives by the end of Fiscal Year (FY) 2024 in order to reinforce the Government's defenses against increasingly sophisticated and persistent threat campaigns. Those campaigns target Federal technology infrastructure, threatening public safety and privacy, damaging the American economy, and weakening trust in Government.

I. OVERVIEW

Every day, the Federal Government executes unique and deeply challenging missions: agencies¹ safeguard our nation's critical infrastructure, conduct scientific research, engage in diplomacy, and provide benefits and services for the American people, among many other public functions. To deliver on these missions effectively, our nation must make intelligent and vigorous use of modern technology and security practices, while avoiding disruption by malicious cyber campaigns.

Successfully modernizing the Federal Government's approach to security requires a Government-wide endeavor. In May of 2021, the President issued Executive Order (EO) 14028, *Improving the Nation's Cybersecurity*,² initiating a sweeping Government-wide effort to ensure that baseline security practices are in place, to migrate the Federal Government to a zero trust architecture, and to realize the security benefits of cloud-based infrastructure while mitigating associated risks.

¹ As used in this memorandum, "agency" has the meaning given in 44 U.S.C. § 3502.

² Exec. Order No. 14028, 86 Fed. Reg. 26633 (2021). <https://www.federalregister.gov/d/2021-10460>

II. EXECUTIVE SUMMARY

In the current threat environment, the Federal Government can no longer depend on conventional perimeter-based defenses to protect critical systems and data. As President Biden stated in EO 14028, “Incremental improvements will not give us the security we need; instead, the Federal Government needs to make bold changes and significant investments in order to defend the vital institutions that underpin the American way of life.”

A transition to a “zero trust” approach to security provides a defensible architecture for this new environment. As described in the Department of Defense Zero Trust Reference Architecture,³ “The foundational tenet of the Zero Trust Model is that no actor, system, network, or service operating outside or within the security perimeter is trusted. Instead, we must verify anything and everything attempting to establish access. It is a dramatic paradigm shift in philosophy of how we secure our infrastructure, networks, and data, from verify once at the perimeter to continual verification of each user, device, application, and transaction.”

This strategy envisions a Federal Government where:

- Federal staff have enterprise-managed accounts, allowing them to access everything they need to do their job while remaining reliably protected from even targeted, sophisticated phishing attacks.
- The devices that Federal staff use to do their jobs are consistently tracked and monitored, and the security posture of those devices is taken into account when granting access to internal resources.
- Agency systems are isolated from each other, and the network traffic flowing between and within them is reliably encrypted.
- Enterprise applications are tested internally and externally, and can be made available to staff securely over the internet.
- Federal security teams and data teams work together to develop data categories and security rules to automatically detect and ultimately block unauthorized access to sensitive information.

This strategy places significant emphasis on stronger enterprise identity and access controls, including multi-factor authentication (MFA). Without secure, enterprise-managed identity systems, adversaries can take over user accounts and gain a foothold in an agency to steal data or launch attacks. This strategy sets a new baseline for access controls across the Government that prioritizes defense against sophisticated phishing, and directs agencies to consolidate identity systems so that protections and monitoring can be consistently applied. Tightening access controls will require agencies to leverage data from different sources to make intelligent decisions, such as analyzing device and user information to assess the security posture of all activity on agency systems.

³ Department of Defense (DoD) Zero Trust Reference Architecture, [https://dodcio.defense.gov/Portals/0/Documents/Library/\(U\)ZT_RA_v1.1\(U\)_Mar21.pdf](https://dodcio.defense.gov/Portals/0/Documents/Library/(U)ZT_RA_v1.1(U)_Mar21.pdf)

A key tenet of a zero trust architecture is that no network is implicitly considered trusted—a principle that may be at odds with some agencies’ current approach to securing networks and associated systems. All traffic must be encrypted and authenticated as soon as practicable. This includes internal traffic, as made clear in EO 14028, which directs that all data must be encrypted while in transit. This strategy focuses agencies on two critical and widely used protocols in the near-term, DNS and HTTP traffic;⁴ in addition, the Cybersecurity and Infrastructure Security Agency (CISA) and the Federal Risk and Authorization Management Program (FedRAMP) will evaluate options for encrypting email in transit.

Further, Federal applications cannot rely on network perimeter protections to guard against unauthorized access. Users should log into applications, rather than networks, and enterprise applications should eventually be able to be used over the public internet. In the near-term, every application should be treated as internet-accessible from a security perspective. As this approach is implemented, agencies will be expected to stop requiring application access be routed through specific networks, consistent with CISA’s zero trust maturity model.⁵

In addition to robust internal testing programs, agencies should scrutinize their applications as our nation’s adversaries do. This requires welcoming external partners and independent perspectives to evaluate the real-world security of agency applications, and a process for coordinated disclosure of vulnerabilities by the general public.

This strategy also calls on Federal data and cybersecurity teams within and across agencies to jointly develop pilot initiatives and Government-wide guidance on categorizing data based on protection needs, ultimately building a foundation to automate security access rules. This collaborative effort will better allow agencies to regulate access based not only on who or what is accessing data, but also on the sensitivity of the data being requested.

Transitioning to a zero trust architecture will not be a quick or easy task for an enterprise as complex and technologically diverse as the Federal Government. The strategy set forth in this memorandum is designed to reduce uncertainty and outline a common path toward implementing EO 14028, by updating and strengthening information security norms throughout the Federal enterprise.

III. ACTIONS

While the concepts behind zero trust architectures are not new, the implications of shifting away from “trusted networks” are new to most enterprises, including many agencies. This process will be a journey for the Federal Government, and there will be learning and adjustments along the way as agencies adapt to new practices and technologies.

Agencies that are further along in their zero trust process should partner with those still beginning by exchanging information, playbooks, and even staff. Agency Chief Financial

⁴ DNS is the internet’s Domain Name System, and in this context refers to the protocol used to look up the internet protocol (IP) address of a given hostname (e.g. whitehouse.gov). HTTP stands for Hypertext Transfer Protocol, and is the primary protocol used to serve web content, as well as other internet data.

⁵ CISA, *Zero Trust Maturity Model*, <https://cisa.gov/publication/zero-trust-maturity-model>

Officers, Chief Acquisition Officers, senior agency officials for privacy, and others in agency leadership should work in partnership with their IT and security leadership to deploy and sustain zero trust capabilities. It is critical that agency leadership and the entire “C-suite” be aligned and committed to overhauling an agency’s security architecture and operations.

Agencies should make use of the rich security features present in cloud infrastructure. This strategy frequently references cloud services, but also addresses on-premise and hybrid systems.

Although this memorandum directs agencies to the highest-value starting points on their path to a zero trust architecture, and describes several shared services which should be prioritized to support a long-term Government-wide effort, this strategy is a starting point, not a comprehensive guide to a fully mature zero trust architecture. In planning and executing their long-term security architecture migration plans, agencies can reference the comprehensive maturity models and reference architectures provided in Appendix A.

This memorandum requires agencies to achieve specific zero trust security goals by the end of Fiscal Year (FY) 2024. These goals are organized using the zero trust maturity model developed by CISA. CISA’s zero trust model describes five complementary areas of effort (pillars) (Identity, Devices, Networks, Applications and Workloads, and Data), with three themes that cut across these areas (Visibility and Analytics, Automation and Orchestration, and Governance).

The strategic goals set forth in this memorandum align with CISA’s five pillars:

1. **Identity:** Agency staff use enterprise-managed identities to access the applications they use in their work. Phishing-resistant MFA protects those personnel from sophisticated online attacks.
2. **Devices:** The Federal Government has a complete inventory of every device it operates and authorizes for Government use, and can prevent, detect, and respond to incidents on those devices.
3. **Networks:** Agencies encrypt all DNS requests and HTTP traffic within their environment, and begin executing a plan to break down their perimeters into isolated environments.
4. **Applications and Workloads:** Agencies treat all applications as internet-connected, routinely subject their applications to rigorous empirical testing, and welcome external vulnerability reports.
5. **Data:** Agencies are on a clear, shared path to deploy protections that make use of thorough data categorization. Agencies are taking advantage of cloud security services to monitor access to their sensitive data, and have implemented enterprise-wide logging and information sharing.

EO 14028 required agencies to develop their own plans for implementing zero trust architecture. **Within 60 days of the date of this memorandum,** agencies must build upon those plans by incorporating the additional requirements identified in this document and submitting to OMB and CISA an implementation plan for FY22-FY24 for OMB concurrence, and a budget

estimate for FY24. Agencies should internally source funding in FY22 and FY23 to achieve priority goals, or seek funding from alternative sources, such as working capital funds or the Technology Modernization Fund.

Agencies will have 30 days from the publication of this memorandum to designate and identify a zero trust strategy implementation lead for their organization. OMB will rely on these designated leads for Government-wide coordination and for engagement on planning and implementation efforts within each organization.

OMB and CISA will work with agencies throughout zero trust implementations to capture best practices, lessons learned, and additional agency guidance on a jointly maintained website at zerotrust.cyber.gov.

A. Identity

Vision

Agency staff use enterprise-managed identities to access the applications they use in their work. Phishing-resistant MFA protects those personnel from sophisticated online attacks.⁶

Actions

1. Agencies must employ centralized identity management systems for agency users that can be integrated into applications and common platforms.
2. Agencies must use strong MFA throughout their enterprise.
 - MFA must be enforced at the application layer, instead of the network layer.
 - For agency staff, contractors, and partners, phishing-resistant MFA is required.
 - For public users, phishing-resistant MFA must be an option.
 - Password policies must not require use of special characters or regular rotation.
3. When authorizing users to access resources, agencies must consider at least one device-level signal alongside identity information about the authenticated user.

1. Enterprise-wide identity systems

The Federal Government must improve its identity systems and access controls. As agencies adopt new infrastructure and applications, they should ensure that information is accessed by the right users, at the right time, and for the right purposes. Doing this well requires two fundamental elements: (1) a holistic view of users, with a strong understanding of their responsibilities and authorities, and (2) an ability to verify the identities of users when they attempt to access systems. Those fundamental elements help agencies establish risk-based access. Doing this effectively requires implementing strong authentication across the enterprise

⁶ In this document, "phishing-resistant" authentication refers to authentication processes designed to detect and prevent disclosure of authentication secrets and outputs to a website or application masquerading as a legitimate system.

and consolidate the means of authenticating to as few agency-managed identity authentication systems as practicable.

Zero trust architectures require metadata about the user to allow agencies to make risk-based decisions at the policy enforcement point. That metadata is maintained, updated, and supplied by systems that manage user identities, keeping the appropriate metadata associated with the correct user even if that user leaves the organization or moves to a new position within it. Such enterprise identity systems integrate with and draw data from external systems, such as those dedicated to human resources, contract management, or personnel security, to gain time-relevant information about the user.

Using centrally managed systems to provide enterprise identity and access management services reduces the burden on agency staff to manage individual accounts and credentials. It also improves agencies' knowledge of user activities, thereby enabling better detection of anomalous behavior, allowing agencies to more uniformly enforce security policies that limit access, as well as quickly detect and take action against anomalous behavior when needed.

Given the importance and advantages of enterprise identity and access management, each Federal agency should support well-designed enterprise identity management systems to perform these functions and integrate it into as many agency applications as possible. Large agencies with many different systems requiring user authentication will only be able to efficiently perform baseline operations, such as promptly disabling the accounts of departing employees, by consolidating authentication. Such consolidation is also critical if large agencies are to implement some of the more sophisticated protections required by this memorandum.

Enterprise identity management must be compatible with common applications and platforms. As a general matter, users should be able to sign in once and then directly access other applications and platforms within their agency's IT infrastructure. Beyond compatibility with common applications, an agency identity management program should facilitate integration among agencies and with externally operated cloud services; the use of modern, open standards often promotes such integration.

It is important to note these decisions are not typically isolated to one agency. It is common practice for agencies to offer services to other agencies. Federated trust relationships between agencies and shared authentication services are opportunities for better integration and coordination.

To promote consistent and auditable identity practices, an agency's enterprise identity systems should also be capable of supporting human authentication through non-graphical user interfaces, such as scripts and command line tools.

2. Multi-factor authentication

Strong authentication is a necessary component of a zero trust architecture, and MFA will be a critical part of the Federal Government's security baseline.

Agencies must integrate and enforce MFA across applications involving authenticated access to Federal systems by agency staff, contractors, and partners.⁷

MFA should be integrated at the application layer, such as through an enterprise identity service as described above, rather than through network authentication (e.g., a virtual private network).

Approaching an application from a particular network must not be considered any less risky than approaching it from the public internet. Accomplishing this goal in an enterprise means progressively de-emphasizing network-level authentication by its users, and eventually removing it entirely. In mature zero trust deployments, users strongly authenticate into applications, not into the underlying networks.

MFA will generally protect against some common methods of gaining unauthorized account access, such as guessing weak passwords or reusing passwords obtained from a data breach. However, many approaches to multi-factor authentication will not protect against sophisticated phishing attacks, which can convincingly spoof official applications and involve dynamic interaction with users. Users can be fooled into providing a one-time code or responding to a security prompt that grants the attacker account access. These attacks can be fully automated and operate cheaply at significant scale.

Fortunately, there are phishing-resistant approaches to MFA that can defend against these attacks. The Federal Government's Personal Identity Verification (PIV) standard is one such approach. The World Wide Web Consortium (W3C)'s open "Web Authentication" standard,⁸ another effective approach, is supported today by nearly every major consumer device and an increasing number of popular cloud services.

Agencies must require their users⁹ to use a phishing-resistant method to access agency-hosted accounts. For routine self-service access by agency staff, contractors, and partners, agency systems must discontinue support for authentication methods that fail to resist phishing, including protocols that register phone numbers for SMS or voice calls, supply one-time codes, or receive push notifications.

This requirement for phishing-resistant methods is necessitated by the reality that enterprise users are among the most valuable targets for phishing. That problem can be mitigated by providing those users with phishing-resistant tokens, including the PIV cards that agency staff and partners are generally issued.

⁷ The term "partners" is meant to include users that are external to the agency, but whose use of agency systems requires a strong form of MFA. For example, this category could include Government contractors submitting financial information. Agencies will need to determine the scope of this category based on their own systems and missions.

⁸ Web Authentication, also known as WebAuthn, was developed as part of the FIDO Alliance's FIDO2 standards, and is now published by the World Wide Web Consortium (W3C) as a free and open standard: <https://www.w3.org/TR/webauthn-2/>

⁹ These users include employees, contractors, and enterprise users, such as a mission or business partners, as described in OMB Memorandum M-19-17. <https://www.whitehouse.gov/wp-content/uploads/2019/05/M-19-17.pdf>

For many agency systems, PIV (including Derived PIV¹⁰) will be the simplest way to support phishing-resistant MFA requirements, and OMB Memorandum M-19-17 requires agencies to use PIV credentials as the “primary” means of authentication to Federal information systems.

However, PIV will not be a practical option for some information systems and situations. Agencies are permitted under current guidance to use phishing-resistant authenticators that do not yet support PIV or Derived PIV (such as FIDO2 and Web Authentication-based authenticators) in order to meet the requirements of this strategy. To the greatest extent possible, agencies should centrally implement support for non-PIV authenticators in their enterprise identity management systems, so that these authenticators are centrally managed and connected to enterprise identities.

Agencies are still expected to maintain exceptional procedures for emergency situations and account recovery processes. By their nature, recovery processes represent a potential bypass of standard authentication protocols, and thus can be a significant threat vector if not mitigated. Agency recovery processes should be designed with the expectation that they are exceptional, and require high-friction methods that are costly for an adversary to overcome, such as in-person verification, live video interaction, or other similar methods.

Privileged Access Management (PAM) solutions that provide ephemeral single-factor credentials for human access to a system should not be used as a general purpose substitute for multi-factor authentication, or for routine single-sign-on access to legacy systems in place of needed modernization of those systems. However, they are still an important tool for improving the security of high privilege systems that are difficult or infeasible to modernize in the near term.

Agencies are encouraged to pursue greater use of passwordless multi-factor authentication as they modernize their authentication systems. However, when passwords are in use, they are a “factor” in multi-factor authentication. If outdated password requirements lead agency staff to reuse passwords from their personal life, store passwords insecurely, or otherwise use weak passwords, adversaries will find it much easier to obtain unauthorized account access—even within a system that uses MFA.

Consistent with the practices outlined in SP 800-63B, agencies must remove password policies that require special characters and regular password rotation from all systems within one year of the issuance of this memorandum. These requirements have long been known to lead to weaker passwords in real-world use and should not be employed by the Federal Government. These policies should be removed by agencies as soon as is practical and should not be contingent on adopting other protections.

¹⁰ NIST Special Publication 800-157, *Guidelines for Derived Personal Identity Verification (PIV) Credentials*, <https://csrc.nist.gov/publications/detail/sp/800-157/final>

Additional technical guidelines that will help accommodate a broad range of multifactor authenticators as Derived PIV Credentials will be published in an upcoming revision to SP 800-157.

This memorandum focuses primarily on the internal enterprise security posture of agencies. However, the security of enterprise and public authentication systems are interconnected. Some Federal systems, such as those that process pre-hire background investigations or the financial information of Government contractors, may be technically public-facing, yet have significant, direct impacts on the operation and security of the Government. In addition, using the same technologies for authentication across both enterprise and public systems fosters interoperability and user familiarity, while improving security across the board.

Systems serving the general public may not yet be able to rely on phishing-resistant authentication alone in providing users access to online services, as some users of online Government services may have limited access to up-to-date devices and security technologies. At the same time, online public services are a major target for phishing attacks and account takeover, and many users will expect Government services to give them tools they can use to protect themselves. To equitably balance security and usability, public-facing Government systems need to offer users more options for authentication.

To that end, public-facing agency systems that support MFA must give users the option of using phishing-resistant authentication within one year of the issuance of this guidance. Meeting this requirement for the general public will mean providing support for Web Authentication-based approaches, such as security keys.¹¹ Agencies may also offer support for authentication using PIV and CAC credentials for agency staff and contractors who are accessing public-facing systems in their personal capacity.

3. User Authorization

In addition to authentication, agencies should ensure their tools can execute certain protocols for authorization. Authorization, a critical aspect of zero trust architecture, is the process of granting an authenticated entity access to resources. Authentication helps ensure that the user accessing a system is who they claim to be; authorization determines what that user has permission to do.

Authorization happens after an authentication event and may be performed by a different set of controls from those that performed authentication. In a zero trust architecture, every request for access should be evaluated to determine whether it is appropriate, which requires the ability to continuously evaluate any active session. If undue risk is identified, mitigations could include requiring reauthentication, limiting access until confirmation that the user requested action is appropriate, or denying access entirely.

Currently, many authorization models in the Federal Government focus on role-based access control (RBAC), which relies on static pre-defined roles that are assigned to users and determine their permissions within an organization. A zero trust architecture should incorporate

¹¹ Agencies should not request information on the make or model of user-supplied authenticators for public-facing systems, to limit unnecessary information collection and to maintain flexibility in user choice of authenticator.

more granularly and dynamically defined permissions, as attribute-based access control (ABAC)¹² is designed to do.

Authorization can be performed at multiple levels. For example, coarse-grained authorization—such as determining who gains initial access to an application—might be performed by tools that rely on ABAC-based approaches, such as those described in NIST SP 800-207. Fine-grained authorization, which determines access to particular data, can be performed within an application itself to grant users varying levels of access based on their role (RBAC).

ABAC and RBAC can be used to allow or deny access by enforcing checks based on the user's identity, the attributes of the resource being accessed, and the environment at access-time. For example, information about the device the user is using (is the device known to the agency? are its patches up-to-date?) provides the basis for a common environment-based check. Analyzing multiple attributes can give an agency higher confidence that the user is permitted to perform a requested action.

Agency authorization systems should work to incorporate at least one device-level signal alongside identity information about the authenticated user when regulating access to enterprise resources.

B. Devices

Vision

Agencies maintain a complete inventory of every device authorized and operated for official business and can prevent, detect, and respond to incidents on those devices.

Actions

1. Agencies must create reliable asset inventories through participation in CISA's Continuous Diagnostics and Mitigation (CDM) program.
 - CISA will design the CDM program to better support a cloud-oriented Federal architecture.
2. Agencies must ensure their Endpoint Detection and Response (EDR) tools meet CISA's technical requirements and are deployed widely.
 - Agencies must work with CISA to identify implementation gaps, coordinate the deployment of EDR tools, and establish information-sharing capabilities, as described in M-22-01.

1. Inventorying assets

¹² NIST defines ABAC as: "An access control method where subject requests to perform operations on objects are granted or denied based on assigned attributes of the subject, assigned attributes of the object, environment conditions, and a set of policies that are specified in terms of those attributes and conditions." See: <https://nvlpubs.nist.gov/nistpubs/specialpublications/NIST.SP.800-162.pdf>

A necessary foundation for any enterprise-wide zero trust architecture is a complete understanding of the devices, users, and systems interacting within an organization. For most enterprises, creating and maintaining a complete inventory over time requires tools that can support the dynamic discovery and cataloging of assets.

CISA operates the Continuous Diagnostics and Mitigation (CDM) program, which aims to help agencies achieve this foundational awareness of their own assets across their enterprise. The CDM program provides a suite of services in support of improved detection and monitoring of agency assets.

As directed by EO 14028, Federal civilian agencies must have formalized their participation in CDM via a memorandum of agreement with DHS. Agencies must create ongoing, reliable, and complete asset inventories, including by leveraging the CDM program.

This is especially practical in cloud environments with rich, granular, and dynamic permission systems. CISA will work toward developing the CDM program to better support a cloud-oriented Federal architecture. For example, CISA may choose to support automated asset discovery using the technical interfaces offered by many commercial cloud infrastructure providers.

2. Government-wide endpoint detection and response

EO 14028 emphasizes the importance of proactive detection of cybersecurity incidents, and the need for Government-wide “hunt” capabilities during incident response. To support the executive order, based on recommendations made by CISA to OMB, OMB issued Memorandum M-22-01,¹³ *Improving Detection of Cybersecurity Vulnerabilities and Incidents on Federal Government Systems Through Endpoint Detection and Response*.

To achieve Government-wide EDR coverage, agencies must ensure that their EDR tools meet CISA’s technical requirements and are deployed and operated across their enterprise. Agencies with robust EDR solutions in place may continue to operate those tools, while agencies that lack them will work with CISA to procure them. To enable Government-wide incident response, agencies must work with CISA to identify implementation gaps, coordinate the deployment of EDR tools, and establish information sharing capabilities, as described in M-22-01.

Agencies should expect to establish procedures and technical facilities to make information reported from their EDR tools available to CISA.

Some specialized systems, such as mainframes and connected devices, may not have compatible EDR tools available. These systems are still at risk of compromise or misuse and may require defenses from other zero trust mechanisms to mitigate risk. Other devices (thin clients, for example) may employ a least-privilege design that specifically constrains general

¹³ Available at <https://www.whitehouse.gov/wp-content/uploads/2021/10/M-22-01.pdf>

purpose computing. Such a design may inhibit the use of common EDR tools but also poses less risk of malicious misuse and is consistent with zero trust principles.

Overall, this approach is intended to maintain a diversity of different EDR tools throughout the Government that can support agencies in differing technological environments, while ensuring a baseline of insight into activity across Federal civilian agencies.

C. Networks

Vision

Agencies encrypt all DNS requests and HTTP traffic within their environment and begin executing a plan to break down their perimeters into isolated environments.

Actions

1. Agencies must resolve DNS queries using encrypted DNS wherever it is technically supported.
 - CISA’s Protective DNS program will support encrypted DNS requests.
2. Agencies must enforce HTTPS for all web and application program interface (API) traffic in their environment.
 - Agencies must work with CISA to “preload” their .gov domains into web browsers as only accessible over HTTPS.
3. CISA will work with FedRAMP to evaluate viable Government-wide solutions for encrypted email in transit and to make resulting recommendations to OMB.
4. Agencies must develop a zero trust architecture plan that describes the agency’s approach to environmental isolation in consultation with CISA and submit it to OMB as part of their zero trust implementation plan.

1. Network visibility and attack surface

As agencies broadly encrypt traffic, it will be critical to balance the depth of their network monitoring against the risks of weak or compromised network inspection devices. Inspecting and analyzing logged network traffic is an important tenet of zero trust architecture. At the same time, a key zero trust principle is assuming that any component can be compromised, including monitoring tools. In addition, as CISA¹⁴ and security researchers¹⁵ have warned, network inspection devices can present security vulnerabilities through weak or incorrect implementation of encryption protocols.

For example, agencies should avoid relying on static cryptographic keys with an overly broad ability to decrypt enterprise-wide traffic, as even a brief compromise of such a key would defeat encryption across the agency. Agencies should make heavy internal use of recent versions

¹⁴ “HTTPS Interception Weakens TLS Security,” <https://us-cert.cisa.gov/ncas/alerts/TA17-075A>

¹⁵ Kumar, Deepak, et al., *The Security Impact of HTTPS Interception*, 26th World Wide Web Conference (Apr. 2017), available at <https://jhalderm.com/pub/papers/interception-ndss17.pdf>

of standard encryption protocols, such as TLS 1.3, that are designed to resist bulk decryption. More generally, agencies should plan for cryptographic agility in their network architectures, in anticipation of continuing to adopt newer versions of TLS and other baseline encryption protocols. In practice, as NIST describes in SP 800-207,¹⁶ there may be places where network traffic cannot or should not be deeply inspected. For example, the risks of weak or compromised network inspection devices can be higher for networks that service a diverse and dynamic set of users, devices, and network destinations, such as those used by agency staff for day-to-day work. In addition, as agencies segment their networks, move away from intranets, and permit access to enterprise services from any network, inspecting traffic in these environments will become less practical and less valuable over time.

In other places, deep traffic inspection may be more valuable and can create less of an increase in attack surface. For example, deep traffic inspection could be more appropriate in application environments that guard sensitive data and have a small number of expected network clients and destinations that can be predicted in advance. In general, when decryption and inspection are performed, agencies should employ technologies whose visibility and privileges are constrained to the absolute least necessary to do their jobs.

Network traffic that is not decrypted can and should still be analyzed using visible or logged metadata, machine learning techniques, and other heuristics for detecting anomalous activity. This is consistent with the Trusted Internet Connection (TIC) initiative, as updated in OMB Memorandum M-19-26, which gives agencies the flexibility to maintain appropriate visibility without needing to perform inline traffic decryption.

2. Encrypting DNS traffic

DNS requests are foundational to the operation of enterprise IT and contain data that should be difficult for attackers to intercept or tamper with.

Like many protocols designed in the early days of the internet, DNS requests have traditionally been unencrypted.¹⁷ This has allowed organizations to monitor DNS within their environments through passive network inspection. Unfortunately, this practice allows adversaries many vantage points within an agency environment from which to perform this monitoring as well.

In recent years, updated standards for encrypting DNS requests have emerged and become widely adopted. Given this evolution, agencies should adjust their DNS architecture and associated monitoring to move closer to a zero trust architecture.

Agencies must resolve DNS queries using encrypted DNS wherever it is technically supported. This means that agency DNS resolvers must support standard encrypted DNS

¹⁶ NIST SP 800-207 at 29, <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf>

¹⁷ DNSSEC does not encrypt DNS data in transit. DNSSEC can be used to verify the integrity of a resolved DNS query, but does not provide confidentiality.

protocols (DNS-over-HTTPS or DNS-over-TLS), and must use them to communicate with upstream DNS resolvers. Agency endpoints must enable encrypted DNS in supporting applications (for example, web browsers) and at the operating system level wherever these features are available.¹⁸ If agencies use custom-developed software to initiate DNS requests, they must implement support for encrypted DNS. Agencies should explicitly configure endpoints to use agency-designated encrypted DNS servers, rather than relying on automatic network discovery.

Agencies can continue to identify and log the contents of encrypted DNS requests by accessing this information at the agency’s designated DNS resolvers. Agencies should include in their zero trust migration plan a description of instances in which they have identified a lack of technical support for encrypted DNS. They should also provide their plans to update operating systems or otherwise ensure support for encrypted DNS enterprise-wide by FY24.

Agencies are already required to have DNS requests routed through CISA-operated infrastructure. To support secure agency DNS traffic, CISA’s Protective DNS offering will support encrypted DNS communication and will scale to accommodate use from agency cloud infrastructure and mobile endpoints.

3. Encrypting HTTP traffic

HTTP is the core protocol used for serving applications to web browsers, whether these applications are public or internal-facing. However, beyond user-visible websites, HTTP is also commonly used for many APIs between servers, mobile applications, and other endpoints.

OMB Memorandum M-15-13 and DHS Binding Operational Directive (BOD) 18-01 currently require agencies to use HTTPS, the encrypted form of HTTP, across all internet-accessible web services and APIs. They do not, however, require the use of HTTPS for traffic that is solely internal. Zero trust architectures—and this strategy—require agencies to encrypt all HTTP traffic, including within their environments.

To ensure they meet that requirement, and to strengthen .gov as a top-level domain, agencies must work with the DotGov program at CISA to “preload” agency-owned .gov domains as HTTPS-only in web browsers. Internet domain names can be “preloaded” in web browsers so that those browsers will only access services using those domain names over HTTPS. There are significant security benefits to enforcing HTTPS client-side and domain-wide, and since 2020, the DotGov program has coordinated with web browsers to automatically preload all newly registered .gov domains.

¹⁸ Windows 11, released in October 2021, supports DNS-over-HTTPS: <https://techcommunity.microsoft.com/t5/networking-blog/windows-insiders-gain-new-dns-over-https-controls/ba-p/2494644>

Many preexisting agency .gov domains have not been preloaded up to this point, however.¹⁹ The most significant barrier to doing so has been the presence of “intranet” websites that use publicly registered .gov domains but do not support HTTPS. As agencies encrypt their internal traffic, this barrier will be removed, and agencies will be able to safely preload their domains without risking breakage.

More generally, the .gov top-level domain has announced an intent to eventually preload the entirety of the .gov domain space as an HTTPS-only zone.²⁰ This change would improve the security and zero trust posture of government institutions at all levels throughout the United States that make use of .gov for their enterprise services. However, agencies must do their part to encrypt internal HTTP traffic to minimize breakage and make this transition possible.

4. Encrypting email traffic

It remains challenging today to easily and reliably encrypt an email all the way between any sender and any recipient. Unlike HTTP and DNS, there is not today a clear path forward for guaranteeing that Federal emails are encrypted in transit, particularly for emails with external parties.²¹

However, email remains a critical method of communication and authentication in the operation of everyday life in the Federal Government. Since emails to, from, and within the Federal Government are sent and received by a tremendous diversity of clients and service providers, any solution will necessarily be based on open standards.

CISA will evaluate the viability of current open standards as Government-wide solutions for encrypted email in transit and make recommendations to OMB to inform future Government-wide actions. As part of its evaluation, CISA should partner with FedRAMP to convene and consult with cloud service providers and other participants in the email ecosystem.

5. Enterprise-wide architecture and isolation strategy

In SP 800-207, NIST describes several approaches to a zero trust architecture (ZTA) for enterprise workflows: enhanced identity governance, logical micro-segmentation, and network-based segmentation.²² Each of these approaches has the same goal: to meaningfully isolate environments, so that an adversary that compromises one application or component cannot easily move laterally within an organization and compromise other distinct environments.

The most appropriate approach may vary from agency to agency, depending on the nature of their existing enterprise IT infrastructure and investments and their overall maturity in certain

¹⁹ Preloading of agency .gov domains was referenced by OMB Memorandum M-15-13 and encouraged in implementation guidance, but was not required at issuance: <https://https.cio.gov/guide/#options-for-hsts-compliance>

²⁰ An Intent to Preload, <https://home.dotgov.gov/2020/6/21/an-intent-to-preload/>

²¹ The most common standard for email transit encryption today, STARTTLS, is “opportunistic,” meaning that an attacker can interfere with the secure connection and cause emails to be sent unencrypted. Such attacks have been observed at scale on the public internet.

²² NIST SP 800-207 at 11-13, available at <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf>

zero trust areas (such as identity access management or cloud network architecture). Agencies may find it makes sense to use a combination of approaches across different application environments and operating components within the organization.

Regardless of the approach selected, agencies must move away from the practice of maintaining a broad enterprise-wide network that allows enhanced visibility or access to many distinct applications and enterprise functions. Accordingly, agencies should choose their zero trust approach early enough to permit them to align that approach with their plans for IT investment.

Based on those considerations, this memorandum requires each agency, in consultation with CISA, to develop a zero trust architecture roadmap that describes how the agency intends to isolate its applications and environments. The agency must include that roadmap in the full zero trust implementation and investment plan required by this memorandum. The roadmap should also describe the agency's operational and security objectives for any enterprise-wide network it may currently operate. In addition, the agency should explain how cloud-based infrastructure will fit into the agency's zero trust architecture. Mature cloud platforms typically feature strong identity- and attribute-based access control and rely on identity governance and virtualized logical isolation of environments. As a result, they are well optimized for zero trust architectures, and agencies are expected to make robust, secure use of cloud-based infrastructure.

D. Applications and Workloads

Vision

Agencies treat all applications as internet-connected, routinely subject their applications to rigorous empirical testing, and welcome external vulnerability reports.

Actions

1. Agencies must operate dedicated application security testing programs.
2. Agencies must utilize high-quality firms specializing in application security for independent third-party evaluation.
 - CISA and GSA will work together to make the services of such firms available for rapid procurement.
3. Agencies must maintain an effective and welcoming public vulnerability disclosure program for their internet-accessible systems.
4. Agencies must identify at least one internal-facing FISMA Moderate application and make it fully operational and accessible over the public internet.
5. CISA and GSA will work together to provide agencies with data about their online applications and other assets.
 - Agencies must provide any non-.gov hostnames they use to CISA and GSA.
6. Agencies should work toward employing immutable workloads when deploying services, especially in cloud-based infrastructure.

1. Application security testing

For Federal applications to withstand sophisticated probing and attack, agencies need to go beyond implementing and documenting security controls. To gain confidence in the security of their systems, agencies must analyze their software and its deployed functionality with a comprehensive and rigorous approach, whether their software is built internally or by a contracted vendor.

Agencies already create a Security Assessment Report (SAR) as part of authorizing their information systems. These SARs should incorporate not just information gathered by automated tools for vulnerability scanning and code analysis of custom-developed software, but also analysis prepared by more time-intensive, specialized, and application-specific methods.

For example, running a scanner on a page with a web form to detect common misconfigurations might be a helpful starting point, but would not provide confidence in the security of that form. More thorough testing would be needed. Such testing could involve, for example, attempting to submit creatively invalid data, and evaluating whether data validation is performed consistently on both the client and server.

Agency system authorization processes must employ both automated analysis tools and manual expert analysis. To understand the depth of security analysis that agencies perform on applications prior to authorization, OMB may at any time ask an agency to produce an application's most recent security assessment. Agencies are expected to continue moving toward continuous monitoring and ongoing authorizations while employing periodic manual security assessments as applications, dependencies, components, and infrastructure evolve. Agencies must prioritize and address vulnerabilities identified in their SAR through these methods.

As directed by EO 14028, NIST has developed guidelines for developer verification of software,²³ which should inform agencies' strategies, methodologies, and standard processes for application testing.

2. Easily available third-party testing

In addition to their own testing programs, agencies must increase their reliance on external perspectives to identify vulnerabilities that internal staff may not identify.

To support agencies in achieving this, within one year of publication of this memorandum, CISA and GSA will collaborate to create a procurement structure for agencies that allows for rapid acquisition of rigorous application-security testing capabilities. As a result of this work, agencies should be able to schedule most work within less than a month (or in high-urgency situations, a few days).

²³ National Institute of Standards and Technology, *Guidelines on Minimum Standards for Developer Verification of Software*, (July 2021), available at <https://www.nist.gov/system/files/documents/2021/07/13/Developer%20Verification%20of%20Software.pdf>

3. Welcoming application vulnerability reports

Public vulnerability disclosure programs, which allow security researchers and other members of the general public to report security issues safely, are used widely across the Federal Government and many private-sector industries. These programs are an invaluable accompaniment to existing internal security programs and operate as a reality check on an organization's online security posture.

To ensure agencies are able to receive vulnerability information from the general public, OMB issued Memorandum M-20-32,²⁴ and CISA published Binding Operational Directive 20-01.²⁵ Those authorities require agencies to publish security contact information, as well as a clear and welcoming vulnerability disclosure policy (VDP).

Consistent with these directions, agencies must welcome external vulnerability reports for their internet-accessible systems by September 2022 and structure reporting channels so that system owners have direct, real-time access to incoming vulnerability reports. To improve internal security and avoid public disclosure of unpatched vulnerabilities, agencies should focus on validating and resolving externally reported vulnerabilities in a responsive manner.

To assist agencies, CISA has released a vulnerability disclosure platform²⁶ that agencies may use to receive and triage vulnerability reports and to engage directly with security researchers. FedRAMP will assist agencies by working with cloud platform providers to clarify that Federal agency customers are permitted to authorize vulnerability testing on customer-operated applications and infrastructure hosted on provider platforms.

4. Safely making applications internet-accessible

Making applications internet-accessible in a safe manner, without relying on a virtual private network (VPN) or other network tunnel, is a major shift for many agencies that will take significant effort to achieve. As with all large-scale IT modernization efforts, its chances of long-term success will be improved by beginning with an agile approach.

To catalyze this work and facilitate early identification of obstacles, each agency must select at least one FISMA Moderate system that requires authentication and is not currently internet-accessible. Then, within a year of the issuance of this memorandum, the agency must take the actions necessary to allow secure, full-featured operation of that system over the internet.

Accomplishing that task will require agencies to put in place minimum viable monitoring infrastructure, denial of service protections, and an enforced access-control policy. While implementing those elements, the agency should integrate this internet-facing system into an enterprise identity management system, as described in the Identity section above. Agencies will

²⁴ OMB Memorandum M-20-32, *Improving Vulnerability Identification, Management, and Remediation* (September 2, 2020), available at <https://www.whitehouse.gov/wp-content/uploads/2020/09/M-20-32.pdf>

²⁵ Binding Operational Directive 20-01, *Develop and Publish a Vulnerability Disclosure Policy*, available at <https://cyber.dhs.gov/bod/20-01/>

²⁶ "Secure the Government," <https://bugcrowd.com/programs/organizations/cisa>

likely find it beneficial to gain confidence in their controls and processes by performing this shift first on a FISMA Low system before attempting to meet the requirement of doing so for a FISMA Moderate system.

5. Discovering internet-accessible applications

To effectively implement a zero trust architecture, an organization must have a complete understanding of its internet-accessible assets, so that it may apply security policies consistently and fully define and accommodate user workflows. In practice, it can be very challenging for a large, decentralized organization to track every asset reliably.

For agencies to maintain a complete understanding of what internet-accessible attack surface they have, they must rely not only on their internal records, but also on external scans of their infrastructure from the internet. CISA will provide data about agencies' internet-accessible assets obtained through public and private sources. This will include performing scans of agencies' IT infrastructure. For example, GSA operates a website scanning service²⁷ that measures a variety of useful properties using open source software collaboratively maintained by CISA and GSA.

CISA and GSA will also consult other authoritative data sources, such as .gov domain registrations and DNS request logs. CISA and GSA have access to rich sources of useful information that could be significantly improved with the cooperation of other agencies. Through its operation of the .gov DNS domain,²⁸ CISA has access to an authoritative and complete list of each agency's registered .gov domains. CISA may not, however, know of an agency's use of non-.gov domain names. GSA has historically tracked use of Federal non-.gov web URLs,²⁹ but agency participation in GSA's efforts has been voluntary and incomplete. To create a complete understanding of Federal use of domain names, within 60 days of the issuance of this memorandum, agencies must begin providing CISA and GSA any non-.gov hostnames used by their internet-accessible information systems. CISA and GSA will work with agencies to define a streamlined and mutually agreeable process for cataloging non-.gov hostnames and related data that minimizes manual effort.

6. Immutable workloads

Mature cloud-based infrastructure typically offers technical interfaces that are well-optimized for fully automated deployment strategies and can support deployment and roll-back practices that confer fundamentally improved security properties (also known as DevSecOps).

Automated, immutable deployments support agency zero trust goals by allowing substantially improved least privilege architectures. When application deployments no longer need manual access and in-place intervention, individual access to servers and other resources can be dramatically constrained and more easily centrally managed and audited. In addition,

²⁷ Guide to the Site Scanning Program, <https://digital.gov/guides/site-scanning/>

²⁸ DotGov Program home page, <https://home.dotgov.gov>

²⁹ Government-Managed Domains Outside the .Gov and .Mil Top Level Domains, <https://search.gov/developer/govt-urls.html>

allowing manual changes to the environment inevitably leads to situations where different instances of production servers are running different patches or software versions, increasing the complexity of future deployments and introducing opportunities for error.

In general, code or infrastructure should be deployed to a cloud environment in a way that technically restricts manual modification. Any changes, such as patches to the operating system or software libraries, along with any changes to application code, should be accomplished through a redeployment of the code, service, or infrastructure. Each instance of the infrastructure can be built in the same way, enabling a consistent, homogenous environment.

Agencies should work toward employing immutable workloads when deploying services, especially in cloud-based infrastructure. Modern software development lifecycle practices, including Continuous Integration/Continuous Deployment (CI/CD) and Infrastructure as Code (IaC) facilitate the creation of reliable, predictable, and scalable applications based on immutable workloads.

Agencies should use CISA's Cloud Security Technical Reference Architecture (TRA) as a guide for migrating third party services from on-premise hosting to cloud infrastructure providers. The Cloud Security TRA describes DevSecOps and its key components, cloud migration scenarios, and centralizing support services (such as configuration/change management) to facilitate cloud-based development.

E. Data

Vision

Agencies are on a clear, shared path to deploy protections that make use of thorough data categorization. Agencies take advantage of cloud security services and tools to discover, classify, and protect their sensitive data, and have implemented enterprise-wide logging and information sharing.

Actions

1. Federal Chief Data Officers and Chief Information Security Officers will create a joint committee to develop a zero trust data security guide for agencies.
2. Agencies must implement initial automation of data categorization and security responses, focusing on tagging and managing access to sensitive documents.
3. Agencies must audit access to any data encrypted at rest in commercial cloud infrastructure.
4. Agencies must work with CISA to implement comprehensive logging and information-sharing capabilities, as described in OMB Memorandum M-21-31.³⁰

³⁰ OMB Memorandum M-21-31, *Improving the Federal Government's Investigative and Remediation Capabilities Related to Cybersecurity Incidents*, available at <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>

1. Federal data security strategy

Developing a comprehensive, accurate approach to categorizing and tagging data will be challenging for many agencies. While agencies have been required to inventory their datasets for some time, a comprehensive zero trust approach to data management requires going beyond what agencies may be accustomed to thinking of as “datasets.”

Achieving this goal will not only require developing protections for the packaged datasets agencies store in databases or publish online, but also grappling with more loosely structured and dispersed data systems (such as email and document collaboration) and intermediate datasets that exist principally to support the maintenance of other primary datasets.

To ensure engagement and progress on tackling this challenge, within 90 days of the issuance of this memorandum, the Federal Chief Data Officer (CDO) Council and the Federal Chief Information Security Officer (CISO) Council will create a joint working group on zero trust data security for agencies, with representatives of both councils and led by OMB. This working group will develop a data security guide for agencies that addresses how existing Federal information categorization schemes can support effective data categorization in a security context. The working group will also support developing enterprise-specific data categories that are not addressed by existing Federal categories.

This working group will identify members who will act as leads, or designate leads within their agencies, to convene a community of practice that can assist agencies in tackling specific areas of focus. The working group will work closely with the Interagency Council on Statistical Policy and consult with other Federal councils and key stakeholders during development of the guide described. Because the technology market supporting enterprise-wide data categorization is still maturing, the working group also will identify and support pilots of emerging approaches among agencies.

2. Automating security responses

As agencies grapple with security events throughout their systems and cloud infrastructure, automation of security monitoring and enforcement will be a practical necessity. This capability is often referred to as Security Orchestration, Automation, and Response (SOAR).

Making this sort of automation work in a large enterprise—measurably improving security and efficiency without causing unacceptable disruption to the daily work of the organization—will require careful tuning, iteration, and sensitivity to business needs. For an automated security system to operate effectively with a mostly hands-off approach, false positives and false negatives must be low.

Successful automation of security responses will require rich data to inform systems for orchestration, as well as permission management. This includes the types of data being protected and who is accessing the data.

Agencies should strive to employ heuristics rooted in machine learning to categorize the data they gather, and to deploy processes that offer early warning or detection of anomalous behavior in as close to real time as possible throughout their enterprise. For example, agencies may benefit from detecting excessive access requests to certain data types, or when accounts associated with agency leadership are accessing a system or category of data they have not previously accessed and would ordinarily not be expected to.

Machine learning models can be opaque and complex to refine. Overseeing and configuring software that uses machine learning requires specialized skillsets that will take time to develop. In the short-term, agencies must identify early candidates for data sensitivity categorization and security automation that do not require machine learning in order to be useful and can be achieved using relatively simple technical approaches, such as scripts or regular expressions. Any automated actions should first be implemented in a “report only” mode, in which agency security teams monitor the performance of their heuristics and the accuracy of their categorizations before enabling any security actions that might impact staff workflow.

To get started, within 120 days of the issuance of this memorandum, agency Chief Data Officers must work with key agency stakeholders to develop a set of initial categorizations for sensitive electronic documents within their enterprise, with the goal of automatically monitoring and potentially restricting the sharing of these documents.³¹ These categorizations are expected to be developed manually and do not need to be complete, but should be broad enough to be useful while being specific enough to be reliably accurate.³² For example, an agency that uses a standard template for procurement-sensitive documents could attempt to detect when this template is in use. An agency could monitor for potentially excessive sharing of such documents, whether that sharing occurs through collaboration tools or through email. Depending on the characteristics of a document and the features of an agency’s collaboration suite, an agency may be able to automate the restriction of permissions for viewing this document.

3. Auditing access to sensitive data in the cloud

EO 14028 directs agencies to use encryption to protect data at rest. Encryption at rest can protect data that is copied while at rest, but does not protect against access by compromised system components that are authorized to decrypt data. Cloud-based infrastructure providers now offer a wide variety of services that can help detect that activity, through cloud-managed encryption and decryption operations, with their own associated logs.

By relying on cloud-operated infrastructure to manage keys and access to decryption operations, agencies can still rely on the trustworthiness of associated audit logs even if their own environment is fully compromised. Leaning on third-party infrastructure to enforce security

³¹ Agencies are encouraged to participate in the NIST NCCoE’s project to examine different approaches to data categorization and the implementation of protections based on those categorizations:

<https://www.nccoe.nist.gov/projects/building-blocks/data-classification>

³² For example, detecting documents containing Social Security Numbers simply by looking for 9 digits in a row is unlikely to be reliably accurate.

constraints takes advantage of cloud security tools to implement a stronger zero trust architecture, while also making for more efficient use of agency resources.

When agencies encrypt data at rest in the cloud, agencies must use key management tools to create a trustworthy audit log that documents attempts to access that data. This can be achieved by using key management tools operated by the cloud provider, or key management tools that are on-premise or otherwise external to the agency-controlled cloud environment.

Keys can be customer-managed or provider-managed. The critical requirement for key management is that, even if an application is compromised and an adversary has the ability to decrypt data managed by that application, any decryption attempts will still be reliably logged by a separate system.

At advanced stages of maturity, agencies should combine audit logs with other sources of event data to employ more sophisticated approaches to security monitoring. For example, agencies could compare the timing of data access to the timing of user-initiated events to identify database accesses that may not have been caused by normal application activity.

4. Timely access to logs

EO 14028 calls for decisive action to improve the Federal Government's ability to investigate and recover from incidents and breaches, whether these incidents occur in agency-owned infrastructure or in cloud infrastructure maintained by a third-party provider.

Pursuant to EO 14028, and relying on recommendations from CISA, OMB issued Memorandum M-21-31, *Improving the Federal Government's Investigative and Remediation Capabilities Related to Cybersecurity Incidents*,³³ to establish requirements for the retention and management of logs in cloud-hosted and agency-operated environments. M-21-31 focuses on ensuring centralized access and visibility for the highest-level security operations center (SOC) of each agency and on increasing information-sharing between agencies to accelerate incident response and investigative efforts.

To help agencies prioritize their efforts, Memorandum M-21-31 establishes a tiered maturity model to guide agencies through the implementation of requirements. This maturity model is designed to help agencies balance the adoption of various requirements for implementation, log categorization, improved SOC operation, and centralized access.

Agencies must reach the first event logging maturity level (EL-1) no later than August 27, 2022, as described in Memorandum M-21-31. Among their first priorities, agencies are expected to implement integrity measures limiting access to and allowing cryptographic verification of logs, as well as logging DNS requests made throughout their environment.

³³ Available at <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>

F. OMB Policy Alignment

Moving to a zero trust architecture involves changes to nearly every aspect of an enterprise's security posture. As a result, this strategy necessarily touches on a large number of enterprise security practices, which can intersect with other existing OMB policies. This section describes how agencies should interpret other OMB memoranda whose requirements relate to the zero trust goals described within this memorandum.

1. *OMB M-21-07 - IPv6 and Zero Trust*

Agencies are undergoing a transition to IPv6, as described in OMB Memorandum M-21-07,³⁴ while at the same time migrating to a zero trust architecture. Agencies should coordinate the implementation of these initiatives when they revisit their enterprise network infrastructure and policies.

M-21-07 is not intended to require commercial shared service providers (e.g., ISPs, CSPs, CDNs) to migrate their internal infrastructures to support IPv6 alone. Instead, agencies should prioritize working with shared services platforms to ensure they provide IPv6 support on the interfaces exposed to system owners and other organizations. More generally, the Federal Government's IPv6 transition should not slow the migration to the cloud or zero trust architectures. Agency IPv6 adoption plans should first focus on technology areas where IPv6 support is already mature, while allowing time for other service and product providers to upgrade their offerings.

2. *OMB M-19-17 - PIV and non-PIV authenticators*

For many agency systems, PIV (including Derived PIV) will be the simplest way to support phishing-resistant MFA requirements, and OMB Memorandum M-19-17³⁵ requires agencies to use PIV credentials as the "primary" means of authentication used for Federal information systems.

However, PIV will not be a practical option for some information systems and situations. Agencies are permitted under current guidance to use phishing-resistant authenticators that do not yet support PIV or Derived PIV (such as FIDO2 and Web Authentication-based authenticators) in order to meet the requirements of this strategy. To the greatest extent possible, agencies should centrally implement support for alternative authenticators in their enterprise identity management systems, so that these authenticators are centrally managed and connected to enterprise identities.

³⁴ M-21-07, *Completing the Transition to Internet Protocol Version 6 (IPv6)*, <https://www.whitehouse.gov/wp-content/uploads/2020/11/M-21-07.pdf>

³⁵ M-19-17, *Enabling Mission Delivery through Improved Identity, Credential, and Access Management*, <https://www.whitehouse.gov/wp-content/uploads/2019/05/M-19-17.pdf>

3. **OMB M-19-26 and OMB M-21-31** – Alternatives to network inspection

Current OMB policies neither require nor prohibit inline decryption of enterprise network traffic. Agencies are expected to balance the depth of visibility they need with the risks presented by broadly trusted network inspection devices.

Network traffic that is not decrypted can and should still be analyzed using visible or logged metadata, machine learning techniques, and other heuristics for detecting anomalous activity. This is consistent with the Trusted Internet Connection (TIC) initiative, as updated in OMB Memorandum M-19-26,³⁶ which gives agencies the flexibility to maintain appropriate visibility without needing to perform inline traffic decryption.

OMB Memorandum M-21-31,³⁷ *Improving the Federal Government’s Investigative and Remediation Capabilities Related to Cybersecurity Incidents*, describes required fields that agencies must log consistently throughout their enterprise, including packet capture logs. M-21-31 does not require full traffic inspection, but specifies fields that should be captured when such inspection is in place. M-21-31 describes this conditional requirement:

If agencies perform full traffic inspection through active proxies, they should log additional available fields as described in Appendix C and can work with CISA to implement these capabilities. If agencies do not perform full traffic inspection, they should log the metadata available to them.

4. **OMB M-15-13** – HTTPS for internal connections

OMB Memorandum M-15-13³⁸ requires agencies to encrypt HTTP traffic that travels over the public internet to or from a Federal system, using HTTPS and HTTP Strict Transport Security (HSTS). M-15-13 specifically exempts internal connections, stating, “[T]he use of HTTPS is encouraged on intranets, but not explicitly required.” An “intranet” is defined as “a computer network that is not directly reachable over the public internet.”

This memorandum expands the scope of M-15-13 to encompass these internal connections. Agencies should apply the guidance contained in OMB’s published compliance FAQ, at <https://https.cio.gov/guide/>, to their internal systems.

³⁶ M-19-26, *Update to the Trusted Internet Connections (TIC) Initiative*, available at <https://www.whitehouse.gov/wp-content/uploads/2019/09/M-19-26.pdf>

³⁷ M-21-31, *Improving the Federal Government’s Investigative and Remediation Capabilities Related to Cybersecurity Incidents*, available at <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>

³⁸ M-15-13, *Policy to Require Secure Connections Across Federal Websites and Web Services*, available at <https://www.whitehouse.gov/sites/whitehouse.gov/files/omb/memoranda/2015/m-15-13.pdf>

Because this memorandum also requires that agencies preload their .gov domains in web browsers, agencies are expected to satisfy the HSTS requirements of M-15-13 through preloading, rather than applying distinct HSTS policies to individual services.

ATTACHMENTS

Appendix A: References
Appendix B: Task Matrix

Appendix A: References

The Federal Government has been preparing for the transition to a zero trust architecture for some time. Several agencies have published architectural models that can be helpful to other agencies:

- **CISA’s Zero Trust Maturity Model** is a high-level overview of zero trust “pillars” that shows how agencies may progress to “Advanced” and “Optimal” states and describes how CISA service-offerings align to these pillars. Available at: <https://www.cisa.gov/publication/zero-trust-maturity-model>.
- **CISA’s Cloud Security Technical Reference Architecture**, co-authored with the United States Digital Service and FedRAMP, provides a more granular reference for secure cloud architectures and migration strategies. Available at: <https://www.cisa.gov/publication/cloud-security-technical-reference-architecture>.
- **NIST’s SP 800-207, Zero Trust Architecture** provides a consensus definition and framework for the key tenets of zero trust architecture, while describing several different approaches to zero trust architecture that organizations with different risk postures and skillsets can adopt. Available at: <https://csrc.nist.gov/publications/detail/sp/800-207/final>.
- The NIST National Cybersecurity Center of Excellence (NCCoE) has initiated **Implementing a Zero Trust Architecture**, a collaboration with industry partners to apply the concepts in NIST SP 800-207 to a conventional enterprise architecture. Available at: <https://www.nccoe.nist.gov/projects/building-blocks/zero-trust-architecture>.
- **GSA’s Zero Trust Architecture Buyer’s Guide** can help agencies identify GSA contract vehicles that offer products and services relevant to agency zero trust implementations. Available at: [https://www.gsa.gov/cdnstatic/Zero%20Trust%20Architecture%20Buyers%20Guide%20v11%2020210610%20\(2\).pdf](https://www.gsa.gov/cdnstatic/Zero%20Trust%20Architecture%20Buyers%20Guide%20v11%2020210610%20(2).pdf).
- **The Department of Defense’s Zero Trust Reference Architecture** comprehensively describes potential security features and architectural controls that the Department plans to execute across its systems. Available at: [https://dodcio.defense.gov/Portals/0/Documents/Library/\(U\)ZT_RA_v1.1\(U\)_Mar21.pdf](https://dodcio.defense.gov/Portals/0/Documents/Library/(U)ZT_RA_v1.1(U)_Mar21.pdf).

Appendix B: Task Matrix

Section	Task	Agency Action Timeline (Deadlines Measured From Date of Issuance of This Memorandum)
General	Agencies must submit to OMB and CISA an implementation plan for FY22-FY24 for OMB concurrence, and a budget estimate for FY23-24.	Within 60 days.
Identity	Agencies must employ centralized identity management systems for agency users that can be integrated into applications and common platforms.	Include in agency implementation plan.
Identity	Agencies must require their users to use a phishing-resistant method to access agency-hosted accounts.	Include in agency implementation plan.
Identity	Public-facing agency systems that support MFA must give users the option of using phishing-resistant authentication.	Within one year.
Identity	Agencies must remove password policies that require special characters and regular password rotation from all systems.	Within one year.
Identity	Agency authorization systems should work to incorporate at least one device-level signal alongside identity information about the authenticated user.	Include in agency implementation plan.
Devices	Agencies must create ongoing, reliable, and complete asset inventories, including by leveraging the CDM program.	Include in agency implementation plan.
Devices	Agencies must ensure their EDR tools meet CISA’s technical requirements and are deployed and operated across their agency.	See M-22-01. ³⁹
Devices	Agencies must work with CISA to identify gaps, coordinate on deployment, and establish information sharing capabilities with CISA, as described in M-22-01.	See M-22-01.
Networks	Agencies must resolve DNS queries using encrypted DNS wherever it is technically supported.	Include in agency implementation plan.
Networks	Agencies must enforce authenticated HTTPS for all production HTTP traffic, including traffic that does not cross the public internet.	Include in agency implementation plan.
Networks	Agencies must work with the DotGov program at CISA to “preload” agency-owned .gov domains as HTTPS-only in web browsers.	Include in agency implementation plan.

³⁹ OMB Memorandum M-22-01, *Improving Detection of Cybersecurity Vulnerabilities and Incidents on Federal Government Systems through Endpoint Detection and Response*, available at <https://www.whitehouse.gov/wp-content/uploads/2021/10/M-22-01.pdf>

Networks	Agencies must develop a zero trust architecture plan that describes how the agency plans to isolate its applications and environments, in consultation with CISA, and include it in the full implementation and investment plan required by this memorandum.	Include in agency implementation plan.
Applications and Workloads	Agency system authorization processes must employ both automated analysis tools and manual expert analysis.	Include in agency implementation plan.
Applications and Workloads	Agencies must welcome external vulnerability reports for their internet-accessible systems.	September 2022, consistent with OMB M-20-32 and BOD 20-01.
Applications and Workloads	Agencies must select at least one FISMA Moderate system that requires authentication and is not currently internet-accessible, and securely allow full-featured operation over the internet.	Within one year.
Applications and Workloads	Agencies must begin providing CISA and GSA any non-.gov hostnames used by their internet-accessible information systems.	Within 60 days.
Applications and Workloads	Agencies should work toward employing immutable workloads when deploying services, especially in cloud-based infrastructure	Include in agency implementation plan.
Data	Agency Chief Data Officers must work with key agency stakeholders to develop a set of initial categorizations for sensitive electronic documents within their enterprise, with the goal of automatically monitoring and potentially restricting how these documents are shared.	Within 120 days.