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

AUTHOR: Office of Management and Budget

I. Overview

The United States Government faces increasingly sophisticated and persistent cyber threat campaigns that target its technology infrastructure, threatening public safety and privacy, damaging the American economy, and weakening trust in Government.

Every day, the Federal Government executes unique and deeply challenging missions: agencies 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*,[1] 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] Exec. Order No. 14028, 86 FR 26633 (2021). https://www.federalregister.gov/d/2021-10460

II. Purpose

In the current threat environment, the Federal Government can no longer depend on perimeter-based defenses to protect critical systems and data. Meeting this challenge will require a major paradigm shift in how Federal agencies approach cybersecurity.

As described in the Department of Defense Zero Trust Reference Architecture,[2] "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 zero trust architecture that:

- Bolsters strong identity practices across Federal agencies;
- Relies on encryption and application testing instead of perimeter security;
- Recognizes every device and resource the Government has;
- Supports intelligent automation of security actions; and
- Enables safe and robust use of cloud services.

This strategy does not attempt to describe or prescribe a fully mature zero trust implementation. Nor does it discourage any agency from going beyond the actions described herein. The purpose of this strategy is to put all Federal agencies on a common roadmap by laying out the initial steps agencies must take to enable their journey toward a highly mature zero trust architecture. This recognizes that each agency is currently at a different state of maturity, and ensures flexibility and agility for implementing required actions over a defined time horizon. The strategy also seeks to achieve efficiencies for common needs by calling for government-wide shared services, where relevant. 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. But 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."

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

III. Goals

EO 14028 directs agencies to focus on meeting key baseline security measures across the government, such as universal logging, multi-factor authentication (MFA), reliable asset inventories, and ubiquitous use of encryption, and to adopt a zero trust architecture.

To do this, the U.S. government's security architecture must avoid implicit trust in devices and networks, assume networks and other components will be compromised, and generally rely on the principle of least privilege.

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 Federal agencies. This will be a journey for the Federal Government, and there will be learning and adjustments along the way as agencies and policies adapt to new practices and technologies.

Agencies that are further along in their zero trust process will need to partner with those still beginning by exchanging information, playbooks, and even staff. Agency chief financial officers, chief acquisition officers, and others in agency leadership will need to work in partnership with their IT and security leadership to build the operational model to deploy and sustain zero trust capabilities.

This strategy encourages agencies to make use of the rich security features present in cloud infrastructure, while ensuring that agency systems are appropriately designed to support secure use of cloud systems. This strategy frequently references cloud services, as agencies are broadly expected to continue increasing their use of cloud infrastructure and associated security services. However, the actions in this strategy also address on-premise and hybrid systems.

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. Comprehensive maturity models and reference architectures are listed in Appendix A, and agencies should use them to plan and execute their long-term security architecture migration plans.

Required Actions

This memorandum requires agencies to achieve specific zero trust security goals by the end of Fiscal Year (FY) 2024. Grouped using the five pillars that underpin the zero trust maturity model of the Cybersecurity and Infrastructure Security Agency (CISA)[3], those goals include:

1. **Identity:** Agency staff use an enterprise-wide identity to access the applications they use in their work. Phishing-resistant MFA protects those personnel from sophisticated online attacks.

[3] CISA, "Zero Trust Maturity Model", https://cisa.gov/publication/zero-trust-maturity-model

2. **Devices:** The Federal Government has a complete inventory of every device it operates and authorizes for Government use, and can detect and respond to incidents on those devices.
3. **Networks:** Agencies encrypt all DNS requests and HTTP traffic within their environment, and begin segmenting networks around their applications. The Federal Government identifies a workable path to encrypting email in transit.
4. **Applications**: Agencies treat all applications as internet-connected, routinely subject their applications to rigorous 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**, Departments and Agencies shall build upon those plans by incorporating the additional requirements identified in this document, and submitting to OMB an implementation plan for FY22-FY24 and a budget estimate for FY23-24. Agencies should re-prioritize funding in FY22 to achieve priority goals, or seek funding from alternative sources, such as agency working capital funds or the Technology Modernization Fund.

Departments and Agencies will have 30 days from the publication of this memorandum to designate and identify a zero trust architecture 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.

A. Identity

Vision

Agency staff use an enterprise-wide identity to access the applications they use in their work. Phishing-resistant MFA protects those personnel from sophisticated online attacks.[4]

Actions

1. Agencies must establish a single sign-on (SSO) service for agency users that can be integrated into applications and common platforms, including cloud services.
2. Agencies must enforce MFA at the application level, using enterprise SSO wherever feasible.
    - For agency staff, contractors, and partners: phishing-resistant MFA is required.
    - For public users: phishing-resistant MFA must be an option.
3. Agencies must adopt secure password policies and check passwords against known-breached data.
    - CISA will make available to agencies one or more services that can check passwords privately, without exposing those passwords.

Well-managed identity systems are at the foundation of a zero trust architecture. Weak access controls allow adversaries to gain a foothold in an organization through account takeover. At the same time, having too many identity systems creates inconsistent security controls and makes it challenging to reliably revoke access across an enterprise. To ensure consistently strong access controls for users, agencies should consolidate identity systems and federate access[5] through their agency and to other agencies as needed.

As Federal agencies broaden their use of MFA, their primary goal is to defend their users not only from credential-theft, but also from automated phishing attacks. This section describes a new baseline for MFA implementation across the Government that prioritizes phishing defense.

That baseline includes updated MFA requirements for public-facing systems that will give more options to the general public. The security of public-facing and internal systems is interconnected, and phishing-resistant MFA becomes more usable and reliable for everyone when it is universally available. However, universal access to public services is of paramount importance, and agencies will retain flexibility to ensure that their digital services can equitably serve their intended users.

1. Enterprise-wide identity

[4] In this document, "phishing-resistant" authentication refers to the definition of "verifier-impersonation resistant" authentication from NIST Special Publication 800-63-3: https://pages.nist.gov/800-63-3/sp800-63b.html#verifimpers

[5] NIST defines "federation" in SP 800-63-3 as "A process that allows the conveyance of identity and authentication information across a set of networked systems." https://pages.nist.gov/800-63-3/sp800-63-3.html#federation For example, federation can allow an agency's staff to use their enterprise identity to sign into another agency's systems, if both agencies support it.

The Federal Government must improve its identity systems and access controls.

As agencies adopt cloud-based infrastructure and applications, they must ensure the same level of strong authentication across various platforms. The more separate account systems an agency operates, the more challenging it is to implement strong authentication across the enterprise, and the higher the burden on agency staff to manage credentials across the various applications they need to use for their jobs.

The simplest way for a Federal agency to address these challenges is to support a single well-designed authentication system, and to integrate it into as many applications as possible throughout the agency. For large agencies with many different systems requiring user authentication, consolidating identity systems will be a practical necessity in order to implement some of the more sophisticated protections required by this memorandum. In addition, several of the other zero trust steps called for in this memorandum functionally require enterprise SSO that can integrate at the application level.

Agencies must employ an identity provider and enterprise-wide single sign-on (SSO) service for agency users that can be integrated into applications and common platforms, and begin decommissioning other identity systems. As a general matter, users should be able to sign in once and then directly access other agency applications and platforms. Consistent with zero trust and risk management principles, agencies can apply additional authentication requirements or monitoring for access to more sensitive applications. Such SSO services should use open standards, such as SAML or OpenID Connect, and should be capable of integration into externally operated cloud services as well as agency-hosted applications. Agencies should review all applications used across their enterprise with the goal of migrating away from any use of non-SSO credentials in favor of the agency's primary SSO service.

Agencies should aim ultimately to use a single identity system that serves all internal users. Agencies may have divisions or components that need to use their own identity systems to meet their own distinct mission needs. When an agency component needs a separate identity system from their parent agency, that component must consolidate around a single identity system. This identity system must be able to support the federation of that component's users throughout other identity systems in their agency, and must accept federated identities from an agency-wide identity system.

2. Multi-factor authentication and resisting phishing

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.

Federal agencies must develop implementation plans to integrate and enforce MFA across applications involving authenticated access by agency staff, contractors, and partners.[6] [7]

MFA should be integrated at the application layer, such as through an enterprise SSO service as described above, rather than through network authentication (e.g., a VPN). Approaching an application from a particular network must not be considered any less risky than approaching it from the public internet. Agencies must steadily de-emphasize network-level authentication by their users, and eventually remove it entirely in their enterprise. In mature zero trust deployments, users log into applications, not into 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, and so will help many agency systems meet this baseline. 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. Any other authentication protocol that meets NIST SP 800-63B's definition of "verifier impersonation-resistant" will also resist the kind of phishing described above.

Agency systems must require internal users to use a phishing-resistant method to access their 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, such as protocols that register phone numbers for SMS or voice calls, supply one-time codes, or receive push notifications.

This requirement for phishing-resistant protocols is necessitated by the reality that enterprise users are among the most valuable targets for phishing, but can be given phishing-resistant tokens, such as PIV cards, and be trained in their use. For many agency systems, PIV or derived PIV will be the simplest way to support this requirement. However, agencies' highest priority should be to rapidly implement a requirement for phishing-resistant verifiers, whether this is PIV or an alternative method, such as WebAuthn.

---

[6] 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 non-Federal registrants of .gov domains, or government contractors submitting financial information. Agencies will need to determine the scope of this category based on their own systems and missions.

[7] For clarity, Privileged Access Management (PAM) solutions may not substitute for multi-factor authentication when authenticating human users to a system.

[8] Web Authentication, also known as WebAuthn, is published as a free and open standard: https://www.w3.org/TR/webauthn-2/

It is expected that emergency situations and human-assisted account recovery processes will create an occasional need for weaker forms of authentication. Any use of non-phishing-resistant authentication must be exceptional and requires CIO approval, sufficient monitoring, and periodic reconsideration.

3. Public-facing authentication

This memorandum focuses primarily on the internal enterprise security posture of Federal agencies. However, the security of enterprise and public authentication are interconnected. Many of the same technologies are used for authentication across both enterprise and public systems, fostering interoperability and user familiarity and improving security across the board.

In addition, 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.

Many systems serving the general public are not yet able to rely on phishing-resistant authentication in providing users access to online services. Some users of online government services may only have access to a landline, and may not have a cell phone or smart phone.

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.** Because most of the general public will not have a PIV or CAC card, agencies will have to meet this requirement by providing support for Web Authentication-based approaches, such as security keys.

4. Using strong password policies

Not all Federal systems use passwords -- but 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, attackers will find it much easier to obtain unauthorized account access -- even within a system that uses MFA.

Agency systems must remove password policies that require special characters and regular password rotation from all systems, whether internal or public-facing. These requirements have long been known to lead to weaker passwords in real-world use[9] and should not be employed by the Federal Government. These policies should be removed as soon as is practical and should not be contingent on adopting other protections described in this document or

[9] "Time to rethink mandatory password changes", https://www.ftc.gov/news-events/blogs/techftc/2016/03/time-rethink-mandatory-password-changes

elsewhere. Agencies should refer to NIST SP 800-63B for guidance on requiring passwords or passphrases of appropriate length and other recommended practices.

CISA will identify or make available to agencies one or more services to privately compare user passwords against known-weak and known-breached data, to help agencies protect against reused stolen credentials.[10] These shared services should be operated such that they cannot discern which passwords are being checked.[11] CISA should be responsive to agencies who wish to use new and emerging services for this function, to keep pace in an evolving security landscape. CISA should also consider using known-breached account information associated with official government email accounts to verify whether breached passwords are currently in use on relevant government systems.

Agency systems must use one of the services identified or made available by CISA to check user passwords against sets of known-weak and known-breached passwords. Since agencies should not store passwords in plaintext, these comparisons can generally only occur when the user enters a password, such as during account creation or login. For particularly sensitive systems, agencies should proactively reset passwords for privileged accounts to ensure that they have been checked.

B. Devices

Vision

The Federal Government has a complete inventory of every device it operates and authorizes for Government work, and can detect and respond to incidents on those devices.

Actions

1. Agencies must participate in CISA's Continuous Diagnostics and Mitigation (CDM) program.
    - CISA will ground the CDM program upon the principle of least privilege, and prioritize effective operation in cloud-based infrastructure.
2. Agencies must ensure that every human-operated enterprise-provisioned device has an agency-chosen endpoint detection and response (EDR) tool.
    - CISA will work with agencies to fill gaps in EDR coverage.
    - Agencies must provide CISA with ongoing access to EDR data.

To enforce a zero trust architecture, agencies must monitor and assess the security posture of all of their authorized devices. As agencies make greater use of cloud services, their assets naturally grow and become more spread out across the internet.

[10] CISA currently performs this practice for one of its own information systems, the DotGov registrar: https://home.dotgov.gov/2018/4/17/increase-security-passwords/

[11] An example of such a privacy-preserving approach: https://blog.cloudflare.com/validating-leaked-passwords-with-k-anonymity/

Agencies must know what they have and where they are vulnerable, whether in-house or in the cloud, in order to successfully monitor and improve the security of their endpoints, servers, and other key technical assets.

1. Inventorying assets

CISA operates the Continuous Diagnostics and Mitigation (CDM) program, which provides a suite of services in support of improved detection and monitoring of agency assets. The CDM program is a foundational element of the Federal Government's approach to situational awareness in a zero trust architecture. As directed by EO 14028, **Federal civilian agencies must formalize their participation in CDM via a memorandum of agreement with DHS.**

As CDM participation grows, it will become even more important to structure CDM itself in accordance with zero trust principles. **CISA must assume that its own monitoring infrastructure could become compromised and adjust the CDM program accordingly**. To do this, CISA must design and adapt CDM to require minimal privileges and avoid the use of privileged software agents wherever possible. CISA must also work with agencies to enforce strict privilege constraints on any CDM tools operating in agency environments.

This is especially practical in cloud environments with rich, granular, and dynamic permission systems. CISA must design the CDM program to natively support Internet of Things (IoT) devices and a cloud-oriented Federal architecture.

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. EO 14028 requires CISA to make recommendations to OMB on implementing an endpoint detection and response (EDR) initiative, and OMB to issue guidance to Federal civilian agencies after receiving those recommendations.

To ensure government-wide EDR coverage, **agencies must ensure strong EDR tools are deployed across their agency.** Agencies with robust EDR tools in place will 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 establish information sharing capabilities with CISA, implemented in accordance with upcoming OMB guidance.**

Agencies should anticipate establishing procedures and technical facilities to make information reported from their EDR tools available to CISA. 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 the Federal civilian government.

C. Networks

Vision

Agencies encrypt all DNS requests and HTTP traffic within their environment, and begin segmenting networks around their applications. The Federal Government identifies a workable path to encrypting email in transit.

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.
    - CISA will work with agencies to "preload" their .gov domains into web browsers as only accessible over HTTPS.
3. CISA will work with FedRAMP to evaluate MTA-STS[12] as a viable government-wide solution for encrypted email and to make resulting recommendations to OMB.
4. Agencies must develop a network segmentation plan in consultation with CISA and submit it to OMB.

A key tenet of a zero trust architecture is that no network is implicitly considered trusted, and accordingly, all traffic is encrypted and authenticated. 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 in the near-term on DNS and HTTP traffic. CISA and FedRAMP will evaluate whether MTA-STS can be a viable, and potentially automatable, path to reliably encrypt email in transit.

As agencies broadly encrypt traffic, they will need to balance the depth of their network monitoring against presenting an excessive attack surface. A key zero trust principle is assuming that any component can be compromised, including monitoring services. Agencies should minimize opportunities for adversaries to gain the ability to view or modify traffic.

For example, agencies should avoid relying on static keys with overly broad ability to decrypt enterprise-wide traffic, as the theft of such a key would defeat encryption across the agency. Agencies should make heavy internal use of recent versions of standard encryption protocols, such as TLS 1.3, that are designed to resist bulk decryption.

This means that in practice, as NIST describes in SP 800-207,[13] there may be places where network traffic cannot be deeply inspected. This network traffic can still be analyzed using

[12] "SMTP MTA Strict Transport Security (MTA-STS)", https://www.rfc-editor.org/rfc/rfc8461.html
[13] NIST SP 800-207, section 5.4, p. 29: https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf

visible metadata, machine learning techniques, and other heuristics for detecting anomalous activity. In places where deep traffic inspection is necessary and worth the attack surface, agencies can employ active proxies whose visibility and privileges are constrained to the least necessary to do their jobs.

1. 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.[14] This has allowed organizations to monitor DNS within their environments through passive network inspection. Unfortunately, this allows adversaries many vantage points within an agency environment to perform this monitoring as well.

In recent years, updated standards for encrypting DNS requests have emerged and become widely adopted. Agencies can now 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 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 (e.g., web browsers), and at the operating system level wherever these features are available.[15] If agencies use custom-developed software to initiate DNS requests, they must implement support for encrypted DNS. 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 are already required to have DNS requests route 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.

2. 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.

Agencies have already been required by OMB Memorandum M-15-13 and Binding Operational Directive (BOD) 18-01 to use HTTPS, the encrypted form of HTTP, across all internet-

[14] 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.
[15] Windows 11 is currently expected to support DNS-over-HTTPS: https://techcommunity.microsoft.com/t5/networking-blog/windows-insiders-gain-new-dns-over-https-controls/ba-p/2494644

accessible web services and APIs. Zero trust architectures require agencies to encrypt internal traffic as well.

Agencies must enforce HTTPS for all production HTTP traffic, including traffic that does not cross the public internet. OMB Memorandum M-15-13 and BOD 18-01 require agencies to enforce HTTPS, but only applied to internet-accessible systems. Agencies must now encrypt all HTTP traffic within their environments.

To ensure this is complete, and to strengthen .gov as a top-level domain, **agencies must "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.

However, many preexisting agency .gov domains have not been preloaded.[16] 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 as part of adopting a zero trust architecture, 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.[17] This change would improve the security and zero trust posture of government institutions throughout the United States that make use of .gov for their enterprise services. However, Federal agencies will need to do their part to encrypt internal HTTP traffic to minimize breakage and make this transition possible.

3. 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.[18]

However, email remains a critical method of communication and authentication in the operation of everyday life in the Federal Government. There is no other widely adopted, open, and interoperable standard for cross-organization communication today. We must make progress in this space. Since emails to, from, and within the Federal Government are sent and received by a

[16] 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
[17] https://home.dotgov.gov/2020/6/21/an-intent-to-preload/
[18] 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. Attacks have been observed at scale on the public internet.

tremendous diversity of clients and service providers, any solution will necessarily be based on open standards.

One promising open internet standard for securing emails in transit is MTA-STS,[19] which allows enterprises to publish policies online that instruct global mail servers to strictly enforce encryption for their emails. MTA-STS is seeing some adoption by major cloud email providers, but is not universally supported among mail relays, and its security policies can be technically challenging for organizations to deploy.

CISA will evaluate the viability of MTA-STS as a government-wide solution for encrypted email 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. In addition, CISA should determine the viability of deploying MTA-STS automatically for .gov domain registrants.[20]

4. Segmenting networks around applications

Agencies should be moving towards an end state where every distinct application they run is in its own separate network environment. Multiple applications may rely on specific shared services for security or other purposes, but should not rely on being co-located within a network with those services and should be prepared to create secure connections between them across untrusted networks.

Agencies must develop an implementation plan to segment their networks around individual applications, in consultation with CISA, and include it in the full implementation and investment plan required by this memorandum. This segmentation plan should describe the agency's strategic approach to transitioning their network architecture, including how the agency will employ network virtualization and automated configuration management to easily replicate network security controls.

[19] "SMTP MTA Strict Transport Security (MTA-STS)", https://www.rfc-editor.org/rfc/rfc8461.html
[20] Security expert Andrew Ayer describes a potential approach to automating MTA-STS for customers of DNS providers: https://www.agwa.name/blog/post/mta_sts_automation

D. Applications

Vision

Agencies treat their applications as internet-connected, routinely subject them 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 such firms available for rapid procurement.
3. Agencies must maintain an effective and welcoming public vulnerability disclosure program.
    - CISA will provide a vulnerability disclosure platform that makes it easy for agency system owners to receive reports directly and engage with security researchers.
4. Agencies must identify at least one internal-facing FISMA Moderate application and make it accessible over the public internet, using enterprise SSO.
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.

Zero trust architectures emphasize putting protections as close as possible to the data and operations being protected. Applications are the front-facing attack surface of federal systems, and as system components they are usually necessarily authorized to have broad data access.

At the same time, agencies cannot rely on network perimeter protections to guard their applications from unauthorized access. In the long-term, consistent with CISA's zero trust maturity model, agencies will be expected to stop requiring that access to applications come from specific networks. This means that, as described in the Identity section above, authentication will be done at the application layer, and applications will generally be accessible over the public internet. **In the near-term, every application should be treated as internet-accessible.**

More generally, to protect applications from attack, agencies need to see their applications as their adversaries see them. This means bringing in external partners and independent perspectives to evaluate the real-world security of agency applications, and welcoming the coordinated disclosure of vulnerabilities by the general public.

To do this well, the government must have a complete understanding of what applications agencies currently have online; this means using the internet and the outside world as a reality check to find systems and vulnerabilities that would otherwise be missed.

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 will need to 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 reflect not only automated tools for code analysis of custom-developed software and vulnerability scanning generally, but more time-intensive, specialized, and application-specific analysis.

For example, running a scanner on a page with a web form to detect common misconfigurations might be helpful, but would not be sufficient to gain confidence in the security of that form. More thorough testing could involve attempting to submit creatively invalid data, or evaluating whether client-side validation logic is also consistently validated on the 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 towards continuous monitoring and ongoing authorizations, while employing periodic manual security assessments as applications 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,[21] which agencies should reference when developing their application testing plans. However, NIST's guidance describes a common baseline for many kinds of applications and does not address specialized testing. Agencies will still need to engage in specialized expert analysis of their applications, beyond evaluating common application issues.

2. *Easily available third-party testing*

In addition to their own testing programs, agencies will need to rely on external perspectives to identify vulnerabilities that internal staff may not consider.

To support agencies in achieving this, **CISA and GSA will collaborate on creating a procurement structure for agencies that allows for rapid acquisition of rigorous application security testing** and whose primary goals should be quality and speed. 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).

[21] National Institute of Standards and Technology, "Guidelines on Minimum Standards for Developer Verification of Software," (July 2021). 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 Federal agencies are able to receive vulnerability information from the general public, OMB issued Memorandum M-20-32,[22] and CISA published Binding Operational Directive 20-01.[23] These actions require agencies to publish security contact information, as well as a clear and welcoming vulnerability disclosure policy (VDP).

Agencies must welcome external vulnerability reports for every online system they operate, and structure reporting channels so that system owners have direct, real-time access to incoming vulnerability reports.

To assist agencies, **CISA has released a vulnerability disclosure platform**[24] that Federal agencies may use to receive and triage vulnerabilities and to engage directly with security researchers.

To avoid any issues that might hinder secure cloud adoption, **FedRAMP will work with cloud platform providers to clarify that Federal agency customers are permitted to authorize vulnerability testing** on customer-operated applications and infrastructure 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) is a major shift for many agencies that will take significant effort to achieve. However, 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 to identify any obstacles early in an agency's zero trust journey, **agencies must select at least one FISMA Moderate system that requires authentication and is not currently internet-accessible, and allow access over the internet**. This will require agencies to create minimum viable monitoring infrastructure and policy enforcement to safely allow internet access. This process should also involve integration with their agency's enterprise-wide single sign-on system, as described in the Identity section above. Agencies will likely find it beneficial to gain confidence in their controls and processes by first performing this shift on a FISMA Low system before meeting the requirement of doing so for a FISMA Moderate system.

---

[22] OMB Memorandum M-20-32, "Improving Vulnerability Identification, Management, and Remediation" (September 2, 2020), https://www.whitehouse.gov/wp-content/uploads/2020/09/M-20-32.pdf

[23] Binding Operational Directive 20-01, "Develop and Publish a Vulnerability Disclosure Policy," https://cyber.dhs.gov/bod/20-01/

[24] "Secure the Government," https://bugcrowd.com/programs/organizations/cisa

5. *Discovering internet-accessible applications*

Zero trust architectures require a complete understanding of an organization's internet-accessible assets, to apply security policies consistently and to 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 the reality of what can be found on the internet.

To assist agencies, **CISA will provide agencies with data about their internet-accessible assets as discovered by CISA and GSA through public and private sources**. CISA and GSA will consult public and commercial data based on internet scans and perform scanning themselves where helpful. CISA and GSA will also consult other authoritative data sources, such as .gov domain registrations and DNS request logs.

Through its operation of the .gov DNS domain,[25] CISA has access to an authoritative and complete list of each agency's registered .gov domains, but cannot learn of the agency's use of non-.gov domain names. GSA operates a website scanning service[26] that measures a variety of useful properties, and relies on open source software collaboratively maintained by CISA and GSA. GSA has also historically tracked use of Federal non-.gov web URLs,[27] but agency participation in GSA's efforts is voluntary and incomplete, and the data is limited to websites.

To assist CISA and GSA, **agencies must provide to CISA and GSA, on an ongoing basis, any non-.gov hostnames used by their internet-accessible information systems.** CISA and GSA will work with agencies to define a streamlined, sustainable, and mutually agreeable process that will meet that requirement while minimizing manual effort among all participants and ensuring long-term data quality.

[25] DotGov Program home page, https://home.dotgov.gov
[26] "Guide to the Site Scanning Program," https://digital.gov/guides/site-scanning/
[27] "Government-Managed Domains Outside the .Gov and .Mil Top Level Domains," https://search.gov/developer/govt-urls.html

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. OMB will work with Federal chief data officers and chief information security officers to develop a zero trust data security strategy and associated community of practice.
2. Agencies must perform some 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.[28]

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."

Agencies must not only develop protections for the packaged datasets they store in databases or publish online, but must also grapple with more loosely structured and dispersed data systems (such as email and document collaboration) and intermediate datasets which exist principally to support the maintenance of other primary datasets.

To ensure engagement and progress on tackling this challenge, the Federal Chief Data Officer (CDO) Council and the Federal Chief Information Security Officer (CISO) Council will create a joint committee on zero trust data security for Federal agencies, chaired by OMB. This committee will develop a data categorization and protection guide for Federal agencies, and oversee a community of practice that can assist agencies in tackling specific areas of focus. This committee will consult with other Federal councils and key stakeholders during development of this guide.

2. Automating security responses

[28] "Improving the Federal Government's Investigative and Remediation Capabilities Related to Cybersecurity Incidents," 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

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 it to be practical for an automated security system to operate with a mostly hands-off approach, false positives and false negatives must be low.

At the same time, to successfully automate security events surrounding data, systems for orchestration and permission management will need rich information on the types of data being protected.

Agencies should strive to employ heuristics rooted in machine learning to detect anomalous behavior or categorize the data they use throughout their enterprise. However, machine learning models can be opaque and complex to debug. Overseeing and configuring software that uses machine learning requires specialized skillsets will take time to develop.

In the short-term, agencies will need to 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 simpler technical approaches, such as scripts or regular expressions. Any automated actions should first be implemented in a "report only" mode, where 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, agencies must leverage their chief data officer to develop a set of initial categorizations for sensitive documents within their enterprise, and automatically monitor and potentially restrict how these documents are shared.[29] 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.[30] For example, an agency which 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 this document when shared via collaboration tools or sent through email. Depending on the characteristics of a document and the features in an agency's collaboration suite, an agency could potentially automate the restriction of permissions around viewing this document.

3. Auditing access to sensitive data in the cloud

EO 14028 calls for 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

[29] 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

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

system components that are authorized to decrypt data. However, cloud-based infrastructure providers now offer a wide variety of services that support cloud-managed encryption and decryption operations, with their own associated logs.

By relying on cloud-operated infrastructure to manage keys and gate 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 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 independently operated key management tools to create a trustworthy audit log of access to that data. This can be achieved by using key management tools operated by the cloud provider, or by key management tools that are on-premise or otherwise external to the agency-controlled cloud environment. In either case, access to key management tools and their audit logs must be isolated from the applications whose activity is being logged. This requirement does not apply to data encrypted in on-premise environments because they do not consistently have third-party components available whose trustworthiness could be relied upon in the event of a total agency compromise.

At advanced stages of maturity, agencies can combine these 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 modeled on recommendations from CISA, OMB issued Memorandum M-21-31, "Improving the Federal Government's Investigative and Remediation Capabilities Related to Cybersecurity Incidents,"[31] to establish requirements for the retention and management of logs in cloud-hosted and agency-operated environments. This memo focuses on ensuring centralized access and visibility for the highest-level security operations center (SOC) of each agency and on increased 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.

[31] "Improving the Federal Government's Investigative and Remediation Capabilities Related to Cybersecurity Incidents," 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

Agencies must reach the first incident logging maturity level (IL1) as described in Memorandum M-21-31. Among their first priorities, agencies are expected to implement log integrity measures to limit access and allow cryptographic verification, and to log DNS requests made throughout their environment.

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.
- **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.
- **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.
- 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.
- **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.
- **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.