Advancing Zero Trust Maturity Throughout the Device Pillar

Executive summary

Continued cyber incidents have called attention to the immense challenges of ensuring effective cybersecurity across the federal government, as with many large enterprises, and demonstrate that "business as usual" approaches are no longer sufficient to defend the nation from cyber threats. The government can no longer depend only on traditional strategies and defenses to protect critical systems and data. [1]

A modernized cybersecurity framework—Zero Trust—integrates visibility from multiple vantage points, makes risk-aware access decisions, and automates detection and response. Implementing this framework places network defenders in a better position to secure sensitive data, systems, applications, and services. [2]

This cybersecurity information sheet (CSI) provides recommendations for maturing devices—the Zero Trust device pillar—to effectively ensure all devices seeking access earn trust based on device metadata and continual checks to determine if the device meets the organization's minimum bar for access. The primary capabilities of the device pillar are:

- identification, inventory, and authentication
- detection of unknown devices and configuration compliance checks of known ones
- device authorization using real time inspections
- remote access protections
- hardware updates and software patches
- device management capabilities
- endpoint detection and response for threat detection and mitigation

This CSI further discusses how these capabilities integrate into a comprehensive Zero Trust (ZT) framework, as described in Embracing a Zero Trust Security Model. [2] National Security System (NSS), Department of Defense (DoD), and Defense Industrial Base (DIB) owners and operators should use this and complementary guidance to understand how to take concrete steps for maturing device security by implementing the outlined capabilities.

Contents

Introduction

Cybersecurity threats are increasing and can originate from a variety of sources – from nation-state actors conducting organized campaigns to individual malicious actors seeking an easy payday. To better secure networks from these threats, networks need to transition from traditional defenses to a Zero Trust (ZT) framework. The ZT security model is best illustrated as seven pillars that together comprise the complete cybersecurity posture. The seven pillars are: User, Device, Network & Environment, Application & Workload, Data, Automation & Orchestration, and Visibility & Analytics. Each pillar requires certain criteria and objectives to achieve ZT enactment.

This cybersecurity information sheet (CSI) focuses on the device pillar and includes recommendations for reaching increasing maturity levels of device pillar capabilities.

Having the ability to identify, authenticate, inventory, authorize, isolate, secure, remediate, and control all devices is essential in a ZT approach. Understanding the health and status of devices informs risk decisions, with real time compliance inspections, continuous risk assessments, and automated remediation informing every access request. [3]

In addition to the more common high-level threats to operating systems and application software, ZT capabilities must defend systems from persistent and hard-to-detect threats against devices. Past examples of low-level, persistent threats include:

- LoJax boot rootkit [4]
- MosiacRegressor firmware implant [5]
- UEFI Secure Boot bypasses BootHole [6] and BlackLotus [7]
- Side channel vulnerabilities such as Spectre, Meltdown, Fallout, ZombieLoad, NetSpectre, Downfall, and Inception
- SSD over-provisioning malware [8]

This ZT device pillar CSI prescribes mechanisms to shield devices from low-level, persistent threats over their entire lifecycle. Adoption of a ZT mindset enables organizations to never assume devices within an established environment are secure or that actors cannot hide from defenses in the OS or applications by delving into hardware and firmware. Implementing mature ZT device pillar capabilities enables organizations to assess devices and respond to risks to critical resources in the environment.

For further background on the ZT concept, refer to Embracing a Zero Trust Security Model. [2] For details on user pillar maturation, refer to Advancing Zero Trust Maturity Throughout the User Pillar. [9]

Audience

This CSI provides guidance primarily intended for NSS, DoD, and the DIB, but may be useful for owners and operators of other systems that might be targeted by sophisticated malicious actors. Guidance for system owners and operators is also available via the National Institute of Standards and Technology (NIST), [10] and the Cybersecurity and Infrastructure Security Agency (CISA). [11] This guidance incorporates the DoD ZT guidance [12] referenced at the end of this document.

Background

The President's Executive Order on Improving the Nation's Cybersecurity (EO 14028) and National Security Memorandum 8 (NSM-8) direct the Federal Civilian Executive Branch (FCEB) agencies and NSS owners and operators to develop and implement plans to adopt a ZT cybersecurity framework. [1] [13] ZT implementation efforts are intended to continually mature cybersecurity protections, responses, and operations over time. Progression of capabilities in each of the seven pillars should be seen as a cycle of continuous improvement based on evaluation and monitoring of threats. [2]

Figure 1 depicts the ZT pillars, including the device pillar. The capabilities and milestones for the device pillar component of the ZT maturity model are described in detail throughout this document. Even though they are depicted separately, it is
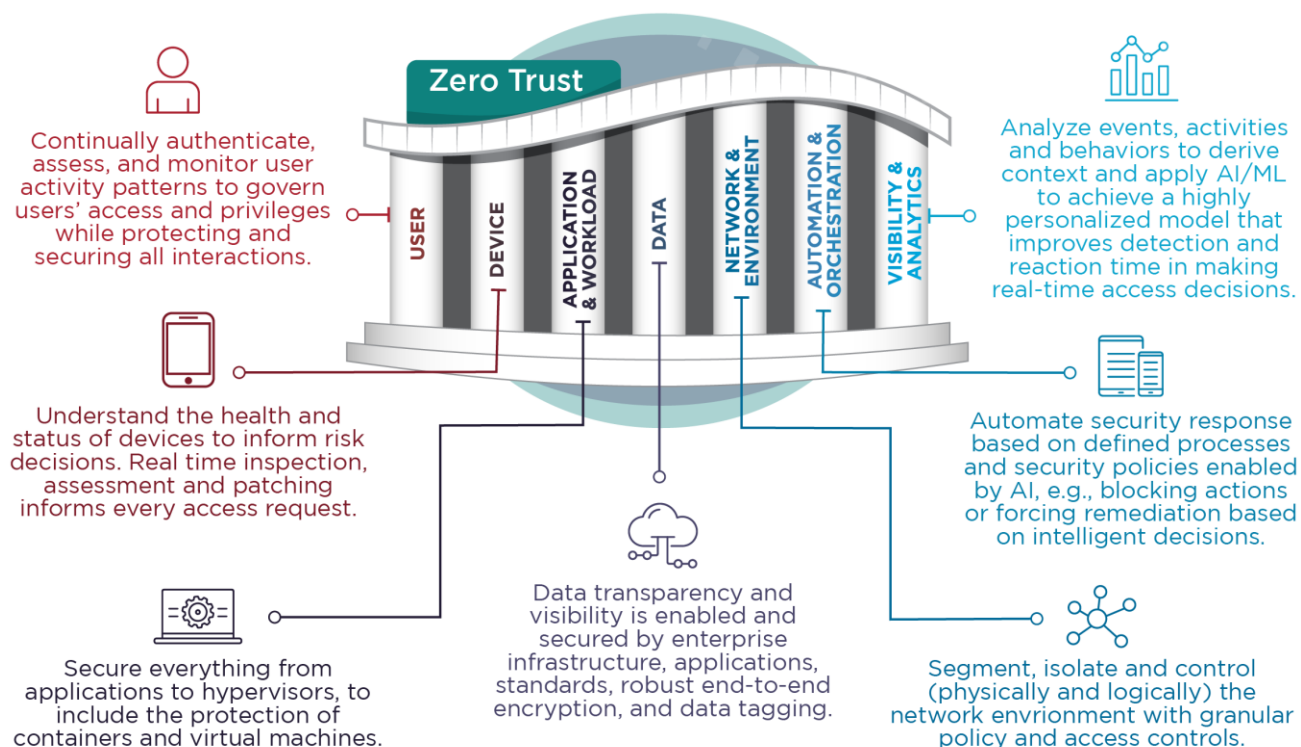

Figure 1: Description of the seven pillars of Zero Trust

important to note that the pillars are not independent; many capabilities in the device pillar depend on or align with capabilities in other pillars, as indicated.

Identity and authentication are based on the user pillar. Devices hosting users are authenticated and authorized to connect to the requested resources based on device attributes. Infrastructure devices are identified and authorized in support of management activities aimed at discovering and responding to threats. Dynamic authentication and authorization decisions are strictly enforced before access is allowed. Recommendations on device connection protocols are included in the network and environment, data, and visibility and analytics pillars. Authentication and remote access are based on the network environment pillar.

Endpoint detection & response (EDR) and extended detection & response (XDR) tools integrate with both the visibility & analytics and the automation & orchestration pillars. EDR/XDR tools enable system administrators to identify, detect, and respond to threats that may be pervasive or present in the environment. Additionally, these security platforms support the necessary analytics that assist with achieving a greater understanding of the performance, behavior, and activities required to improve detection of anomalous behavior to make real time changes in security policies and access decisions.

Device pillar

The device pillar is a foundational component of ZT to ensure devices within an environment, and devices connected to or attempting to connect to resources, are located, enumerated, authenticated, and assessed. Devices are subsequently permitted or denied access — based on a dynamic risk calculation — to specific objects or data. A device is only authorized access if it is "compliant" (meets the environment's security conditions specified by policy). Devices determined to be non-compliant may be denied access or granted limited access.

Each of the following key device pillar capabilities has associated maturity levels:

- **Device Inventory:** Creating device inventory management systems and maintaining real time device inventories. Maintaining a trusted inventory list by enrolling all devices authorized to access the network once they are properly evaluated enables establishing a deny-by-default access policy for devices.

- **Device Detection and Compliance:** Detecting devices as they connect to the network and ensuring compliance with device policies specific to the device function and current risk posture.

- **Device Authorization with Real Time Inspection:** Establishing and utilizing policies to deny devices access to digital resources by default and explicitly allowing access based on compliance, function, and measured risk. Continuous monitoring and behavior analysis enables faster remediation of a broader class of security threats.

- **Remote Access Protection:** Creating policies to allow authenticated and authorized users and devices to access resources from remote locations.

- **Automated Vulnerability and Patch Management**: Identifying the hardware, firmware, and software versions along with their patch levels on devices, correlating them with support information and known vulnerabilities, and upgrading and patching the systems to minimize known risks.

- **Centralized Device Management:** Establishing tooling to manage, secure, and deploy security configurations and applications for computers and mobile devices. In particular, remotely managing and enforcing security policies on organization issued devices.

- **Endpoint Threat Detection and Response:** Implementing tooling to monitor, detect, and remediate malicious activity on devices, integrating with network-wide visibility and defense orchestration capabilities.

As capabilities mature and additional capabilities are deployed, enterprises advance through basic, intermediate, and advanced maturity phases and are more able to operate according to ZT principles.
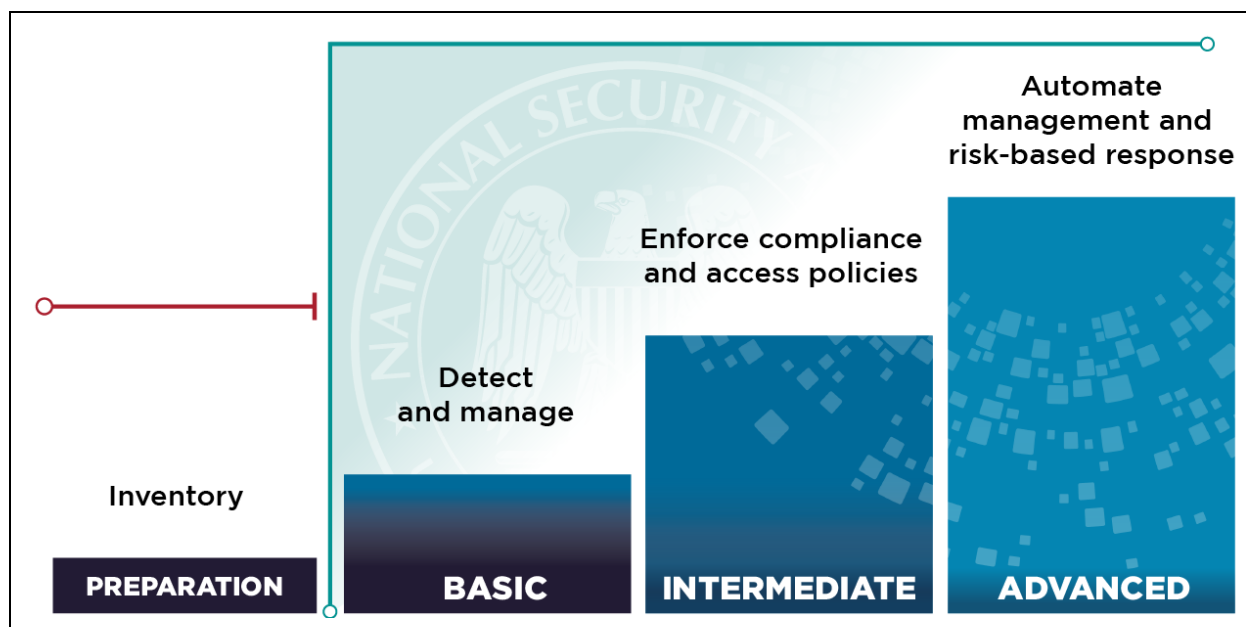

Figure 2: Zero Trust device pillar maturity

Device inventory

Knowing what is in an organization's environment is a foundation to establishing trust in the environment. A device inventory lists what devices are known and expected in the environment. The device inventory can then be used as the basis for starting to establish trust in a device.

A device inventory must capture device existence, usage, and risks. All devices that communicate in an environment require a unique identity and authentication as non-person entities (NPE). Device usage can vary – examples include devices leveraging session access protocols, resource devices hosting or providing network for applications, or devices running embedded services. Enterprises must understand their use of devices and that there is a difference in the way these devices present cybersecurity risks.

The first step to securing the device pillar in a "deny all" by default environment is done by establishing a complete inventory of registered devices that are allowed to access enterprise resources. In some cases, inventory solutions can collect hardware and software information, including versions, patch levels, and installed applications, which are important in establishing security baselines, application allowlisting, and situational awareness across all inventoried devices. Dynamic inventories may include both managed and unmanaged devices that have been granted authorized access to enterprise resources. As maturity increases, dynamic inventories are updated in real time.

Devices may be added or removed from an inventory over time. The action of modifying an inventory requires establishing enterprise policies governing:

- **Procurement:** Identify criteria governing device purchases. Device Authorization – discussed later in this document – may involve the need for specific Trusted Platform Module (TPM) certificates, firmware configuration, or component part revisions. Vendors may list multiple variants or configurations of the same device, but only some may have the necessary components and capabilities.

- **Acceptance Testing:** NIST SP 800-161 calls for enterprises to adopt acceptance testing as a mechanism to audit supply chain integrity. Software Bill of Materials (SBOM), Reference Integrity Manifest (RIM), and TPM Platform Certificate provide artifacts that establish an auditable chain of custody from the production factory to the receiving organization. [14]

- **Deprovisioning:** Devices may store protected data within components other than the storage drive. Plan to securely erase storage media, factory reset firmware, securely erase TPM NVRAM memory, reset Baseboard Management

Controller (BMC) configurations, remove UEFI Secure Boot modifications, and clean up other organization-specific customizations before retiring a device. Inventory should support status records necessary to ensure safe and secure deprovisioning.

Table 1: Device inventory maturity

Preparation	Basic	Intermediate	Advanced
Organizations create an inventory of existing known devices. The inventory is primarily manual and may be based on multiple partial inventories from disparate systems.	Organizations have a complete list of devices in separate inventories. Planning for machine identification and authentication using NPE Public Key Infrastructure (PKI) certificates has started. The organization has identified specific capabilities that must be present on newly acquired assets.	Organizations have a complete list of devices with standardized device attributes and version information. Machine identification and authentication using NPE certificates and a "deny all, allow by exception" approach is mostly implemented. The organization has identified specific make, model, and revisions of devices eligible for new acquisitions. Automation has begun to maintain the device list and bring together disparate inventories.	Organizations have a complete inventory of all devices updated in real time using NPE certificates, enabling only approved devices to be allowed with all others denied by default. An organization acceptance process checks all newly acquired devices and a deprovisioning process sanitizes all devices retired from use.

Device detection and compliance

Networks have many uses and are often intended to be dynamic and adjust to changing uses. Devices entering or leaving the network is part of the expected changes to the network, along with the state of devices changing. Detecting devices and their compliance related to an expected baseline enables managing of the network environment and deciding whether to grant access to devices.

Detection of devices within an environment is achieved through various protocols and solutions. The organization must establish device connection policies that assess device configurations and ensure devices comply with policies established per network and organizational policy. Non-compliant devices represent an unacceptable risk to the organization and should not have access to enterprise resources. For example, one critical area that device configurations affect is the encryption settings a device will use for its communications. In this example, non-compliant configurations could allow the use of obsolete encryption, enabling malicious actors to hijack communications to steal sensitive data, install malware, and other activities. Actions and policies for non-compliant or unknown devices must consider risk posture allowance including ensuring logging, analytics, automated responses, and orchestration.

Organizations must periodically reevaluate compliance policies. Threats to devices may necessitate changes to hardware configuration, firmware version, boot executables, or other device properties over time. Some device vulnerability mitigations may impart a performance impact that requires organizations to balance risk exposure and device performance against organizational objectives.

Table 2: Device detection and compliance maturity

Preparation	Basic	Intermediate	Advanced
Organizations employ asset management systems for user devices to report on compliance with baseline configurations.	Organizations use asset management systems for different types of devices to report compliance. Compliance violations should be logged for later remediation if appropriate.	Organizations have established a minimum selection of compliance attributes and acceptable values. Organizations use asset management systems to track device configurations and check for compliance when devices request to connect to the network, denying access for non-compliance.	Organizations track configurations on all devices, check for compliance continuously, and automatically remediate non-compliance when identified. When remediation is not feasible, the organization uses established, risk-based criteria specific to the device function and capabilities in determining whether to allow access and how much.

Device authorization with real time inspection

Managing a ZT architecture means actively checking that devices in the environment should be trusted for access to critical resources. Authorizing those access requests should be based on current checks that the devices should be trusted for access—not just based on a history of being granted access previously.

Making proper authorization decisions requires the most up-to-date information on which to assess the risk of granting access to data or resources, using information from the Device Detection and Compliance capability combined with real time inspection of additional compliance information as needed. For example, real time inspection may compare current device properties against those from the recorded inventory, examine the device's current patch status, or look for unexpected credentials or applications on the device. Authorization with real time inspection provides continual status updates of a device and its behavior to the decision points making the access decisions. Organizations should establish continual authentication policies to ensure re-authentication of devices when new data or resource accesses are initiated. Each device must be associated with both its current and expected state.

Table 3: Device authorization with real time inspection maturity

Preparation	Basic	Intermediate	Advanced
None at this level.	Organizations provision devices with a unique identifier and are individually authorized.	Organizations use device tooling (e.g., NextGen AV, Application Control, File Integrity Monitoring (FIM), EDR) integration to better understand the risk posture of a device. Access decisions leverage the risk posture and account for device integrity, authentication, and encryption.	Organizations integrate device activity data into risk decisions as well for real time risk assessment of device behavior. All access requests are continuously vetted prior to allowing access to any enterprise or cloud assets.

Remote access protection

When organizations allow remote and hybrid work environments, it is imperative that they authenticate and monitor all internal and external devices that request access to

protected resources. Challenges organizations faced using the conventional architecture was that the user's credentials alone were treated as adequate to grant access to network resources. In a mature ZT architecture, all devices, internal and external, are continually authenticated and monitored.

In particular, organizations should assume a remote user's environment is hostile and that all traffic is being monitored and potentially modified by threat actors, so additional scrutiny of those devices and their access requests is needed. If remote access is authorized, cybersecurity policies, standards, and procedures should include specific policy guidance for required device attributes. Creating a least privilege baseline is critical and should be included for this activity. A thorough authentication, authorization, risk assessment, and determination of acceptable risk must be conducted prior to allowing remote access by all devices.

Organizations should audit existing device access processes and tooling to set a least privilege baseline. Remote access requirements also cover basic bring your own device (BYOD) and Internet of things (IoT) access. They should use the enterprise identity provider (IdP) and only be granted access to approved applications and services when using the acceptable set of device attributes. To accomplish this, BYOD domains may be best governed according to ZT principles utilizing a mobile device management (MDM) tool. Organizations with BYOD environments should look for MDM solutions with separate enrollment policies for employees who want to use their personal devices. [15] [16]

The following table shows remote access maturation from basic to advanced:

Table 4: Remote access protection maturity

Preparation	Basic	Intermediate	Advanced
None at this level.	Organizations employ dynamic access policies with implicit denials, explicit approvals, and centralized management solutions for all remote devices. Control device access to protected	Organizations use centralized management systems to track remote device configurations and check for compliance when devices request to access resources.	All protected services require dynamic access decisions. Automatically remediate non-compliance when identified.

	resources and report compliance.		

Automated vulnerability and patch management

Allowable devices must maintain security updates and patches, otherwise they add significant known mitigatable risks to the network. Having known vulnerabilities does not build trust in devices, instead it should decrease trust. A Z architecture should mitigate risks as much as possible, especially known vulnerabilities that can be patched. A 2023 patch management study found large companies manage at least 2,900 applications across all devices, but more than half of them are not up to date with the latest patches. [17] Automating vulnerability and patch management is critical to protecting resources by defining a security baseline and denying access if this baseline is not met. Threat actors constantly probe for known vulnerabilities – 'low-hanging fruit' that provide an entry route into the targeted environment. Keeping firmware, software, and operating systems up to date reduces the likelihood of being breached. Patches and updates should be tested before implementation to ensure environment stability and that applications continue to function. However, they should be prioritized and tested in a timely manner so that devices are not left vulnerable longer than necessary.

This capability is a special case of the Device Detection and Compliance capability combined with the Centralized Device Management capability to address critical known vulnerabilities since they present a high risk to organizations' devices. In many cases, centralized device management solutions can automate vulnerability identification based on known versions and vulnerabilities and can deploy the necessary patches and updates.

Organizations must maintain awareness of firmware patches below the software layer. These patches may not be delivered via OS patch managers or other automated patching solutions. Some patches may come from the system vendor, while others may be specific to an individual component manufacturer (e.g., SSD firmware provided by the storage vendor – not the system vendor). There are two general realms of device-specific patches:

1. Fixed System firmware: System vendors collaborate with soldered component vendors to deliver patches to customers. CPU microcode and NIC (network interface card) firmware is usually shared by the device's manufacturer.
2. Component firmware: Most frequently applies to components with standardized connectors such as storage drives or graphics processors. Individual component vendors provide firmware updates for their specific products.

Table 5: Automated vulnerability and patch management maturity

Preparation	Basic	Intermediate	Advanced
Organizations track vulnerabilities and apply patches manually.	Organizations use automated feeds to become aware of patches. Patches are manually tested before deployment. All unsupported devices, including any unsupported hardware components or software are identified with plans for their upgrade or retirement.	Organizations use automated tests to check patches for reliability. Once tests are complete, patches are manually approved for automated deployment to all applicable devices according to a schedule intended minimize exposure. All unsupported devices have been removed from the network. Manual or automated (if available) processes to maintain firmware are instituted.	Organizations use automated feeds to trigger patch download and initial automated testing, followed by automated rollout sequencing with automated log and performance analysis to ensure reliability for continued rollout. Any devices that become unsupported are automatically flagged for possible quarantine and upgrade or removal. Organizations also leverage automated asset acceptance testing knowledge to carry out component updates on specific devices when appropriate.

Centralized device management

Knowing that devices are configured securely and managed properly helps build trust in them to then trust them with access to resources. Using centralized device management tools allow the Information Technology team to manage, secure, and deploy corporate resources and applications on any device from a single console. It grants organizations the ability to centrally manage endpoint devices from a single location. Additionally, it provides management with a single view of users that utilize more than one device and assists with retrieving workplace analytics regarding them. [18] It can also improve workplace productivity by continuously providing application and content access to devices. These tools provide a method for organizations to manage all devices from one central location, regardless of what platform they function in. These centralized device management tools are often called Unified Endpoint Management (UEM) solutions for traditional IT devices and Mobile Device Management (MDM) solutions for mobile devices.

Table 6: Centralized device management maturity

Preparation	Basic	Intermediate	Advanced
None at this level.	Organizations employ centralized device management solutions to confirm device compliance status for user devices and report if a device's compliance meets minimum standards.	Organizations have started integrating centralized device management (both UEM and MDM solutions as needed) with inventory capabilities for automated, dynamic inventory of devices combined with device management for compliance. Organizations check the integrity of devices by collecting device integrity values from the TPM and similar device integrity mechanisms.	Organizations inventory all devices via an automated management solution for all services. Security vulnerabilities are identified and patched or mitigated automatically by the device management solutions. Policy is enforced through IT remote management of issued mobile devices. Device integrity values are collected and compared to Software Bill of Materials (SBOM) and Records Information Management (RIM) relevant to the device.

Endpoint threat detection and response

Endpoint threat detection is an essential component of ZT for the device pillar since malicious activity is assumed to be happening at any time. Devices are expected to detect those activities and actively respond to them to contain any damage and remediate the issue. Devices are not inherently trusted, so local threat detection capabilities on the device are used as one capability to build trust that the device is secure. Endpoint threat detection includes local malware prevention solutions, such as antivirus protections, along with other solutions that detect malicious or anomalous behaviors on the device. Combining threat detection with response options enables the device to protect itself from malicious threats. Additionally, reporting of detections and anomalies to centralized visibility and orchestration capabilities (discussed in later pillars) enables awareness by network defenders and appropriate system or network-

wide responses to sophisticated threats. Endpoint threat detection and response often utilizes abilities of Endpoint Detection and Response (EDR) or Extended Detection and Response (XDR) products.

EDR capabilities build upon prior generation Endpoint Security Systems (ESS) by enabling integration of endpoint knowledge with Security Information and Event Management (SIEM) platforms, Security Orchestration, Automation, and Response (SOAR) platforms, incident response activities, and other ZT concepts. XDR platforms further increase visibility and detection of cross-device threats by enabling the correlation of artifacts from endpoints that differ in design, location, or hardware. Correlation of disparate endpoint and environment information is a key maturity measurement associated with advanced ZT, and implementation of XDR will enable organizations to account for activity beyond traditional endpoints.

XDR implementation activities are closely related to SIEM/SOAR capabilities within the Visibility & Analytics and Automation & Orchestration ZT pillars and may include features that support, enhance, or streamline the deployment of other ZT concepts. Robust EDR/XDR deployment can also provide enhancements to:

- Endpoint coverage (visibility & response) across differing device hardware and software.

- Standardization of management interfaces, logging formats, APIs, and endpoint security software footprints.

- Integration of EDR/XDR with activities that reside in other ZT pillars, such as Visibility & Analytics, Automation & Orchestration, and Application & Workload, and can have compounding effects on achieving higher maturity levels.

Other considerations for EDR/XDR implementation:

- EDR platforms benefit from integration with Threat Intelligence and Threat Reputation providers. Endpoint connectivity should be evaluated to the greatest extent possible when assessing the performance of a solution stack.

- Evaluation of a solution stack should take other ZT pillar capability requirements into consideration since EDR/XDR will have direct correlation to the achievement of other ZT pillar capabilities.

- EDR/XDR solutions have varying levels of protection features that require suitability evaluation for each environment. Ensure the solution provides detection, response, or remediation that corresponds with incident response activity requirements and expectations.

Table 7: Endpoint threat detection and response maturity

Preparation	Basic	Intermediate	Advanced
Organizations utilize anti-malware solutions and endpoint auditing services to support manual remediation.	Organizations use EDR solutions to protect, monitor, and respond to malicious and anomalous activities. Organizations prepare to integrate Comply to Connect (C2C) capabilities for expanded device and user checks prior to allowing access. NextGen AV tooling covers maximum number of services/applications.	Organizations utilize XDR solutions to protect, monitor, and respond to malicious and anomalous activities across device types. Integrations with cross-pillar capabilities have been identified and prioritized based on risks. The riskiest integration points are identified and integrated with XDR. Basic alerting sends analytics from XDR stack to the SIEM.	Organizations have completed integrating XDR solutions at all integration points, expanding coverage to fullest capacity. Exceptions are tracked and managed using a risk-based methodical approach. Extended analytics enabling ZT advanced functionalities are integrated into the SIEM and other appropriate solutions.

Summary of guidance

The information presented here is not a standardized solution that fits all organizations, but rather suggestions and considerations for implementing ZT concepts for devices. Discovering and defining the organization's mission and identifying the supporting assets that need to be secured will help build a clearer picture of the as-is architecture which can be compared against the recommendations in this pillar along with the other ZT pillar CSIs. This comparison will help all stakeholders to identify organizational risks and gaps and ultimately inform them on what a mature ZT architecture will look like for their organization. Each organization will need to evaluate their individual requirements to determine a suitable solution. The goal is to develop ZT roadmap strategies that align with the organization's ZT goals. The following guidance are the key ideas for implementing the ZT device pillar:

- Detect and identify devices within or connecting to the environment.
- Authenticate, and continually re-authenticate, devices.

- Use automated solutions to manage device configurations, vulnerabilities, and patches.
- Maintain a dynamic authorization list with policies and procedures in place for denied devices.
- Conduct risk-based assessments to determine access for all devices.
- Enforce more stringent access policies for remote access due to the higher risk environment.
- Monitor endpoints for signs of threat activities, incorporating endpoint monitoring and responses into network-wide security capabilities.

Further guidance

NSA is assisting DoD customers in piloting ZT architectures, coordinating activities with existing NSS and DoD programs, and developing additional ZT guidance to support system developers through the challenges of integrating ZT within NSS, DoD, and the DIB environments. Upcoming additional guidance will help organize, guide, and simplify incorporating ZT principles and designs into enterprise networks.

Works cited

[1]   The White House (2021), Executive Order 14028: Improving the Nation's Cybersecurity. https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/

[2]   NSA (2021), Embracing a Zero Trust Security Model. https:// https://media.defense.gov/2021/Feb/25/2002588479/-1/-1/0/CSI_EMBRACING_ZT_SECURITY_MODEL_UOO115131-21.PDF

[3]   DoD (2022), DoD Zero Trust Strategy. https://dodcio.defense.gov/Portals/0/Documents/Library/DoD-ZTStrategy.pdf

[4]   Ars Technica (2018), First UEFI malware discovered in wild is laptop security software hijacked by Russians. https://arstechnica.com/information-technology/2018/10/first-uefi-malware-discovered-in-wild-is-laptop-security-software-hijacked-by-russians/

[5]   Bleeping Computer (2020), MosaicRegressor: Second-ever UEFI rootkit found in the wild. https://www.bleepingcomputer.com/news/security/mosaicregressor-second-ever-uefi-rootkit-found-in-the-wild/

[6]   Eclypsium (2020), There's a Hole in the Boot. https://eclypsium.com/blog/theres-a-hole-in-the-boot/

[7]   ESET Research (2023), BlackLotus UEFI bootkit: Myth confirmed. https://www.welivesecurity.com/2023/03/01/blacklotus-uefi-bootkit-myth-confirmed/

[8]   Tom's Hardware (2021), New Malware Uses SSD Over-Provisioning to Bypass Security Measures. https://www.tomshardware.com/news/ssd-over-provisioning-vulnerability

[9]   NSA (2023), Advancing Zero Trust Maturity Throughout the User Pillar. https://media.defense.gov/2023/Mar/14/2003178390/-1/-1/0/CSI_Zero_Trust_User_Pillar_v1.1.PDF

[10] NIST (2020), Special Publication 800-207: Zero Trust Architecture. https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf

[11] Cybersecurity and Infrastructure Security Agency (2023), Zero Trust Maturity Model Version 2.0. https://www.cisa.gov/sites/default/files/2023-04/zero_trust_maturity_model_v2_508.pdf

[12] DoD (2022), Zero Trust Reference Architecture Version 2.0. https://dodcio.defense.gov/Portals/0/Documents/Library/(U)ZT_RA_v2.0(U)_Sep22.pdf

[13] The White House (2022), National Security Memorandum 8: Improving the Cybersecurity of National Security, Department of Defense, and Intelligence Community Systems. https://www.whitehouse.gov/briefing-room/presidential-actions/2022/01/19/memorandum-on-improving-the-cybersecurity-of-national-security-department-of-defense-and-intelligence-community-systems/

[14] NIST (2022), NIST Special Publication 800-161r1: Cybersecurity Supply Chain Risk Management Practices for Systems and Organizations. https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-161r1.pdf

[15] NIST (2023), Special Publication 1800-22: Mobile Device Security: Bring Your Own Device (BYOD). https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1800-22.pdf

[16] NIST (2019), NISTIR 8228 Considerations for Managing Internet of Things (IoT) Cybersecurity and Privacy Risks. https://nvlpubs.nist.gov/nistpubs/ir/2019/NIST.IR.8228.pdf

[17] Adaptiva (2023), 2023 Report: The State of Patch Management in the Digital Workplace. https://adaptiva.com/resources/report/state-of-patch-management

[18] Computerworld (2021), What is UEM? Unified endpoint management explained. https://www.computerworld.com/article/3625231/what-is-uem-unified-endpoint-management-explained.html

Disclaimer of endorsement

The information and opinions contained in this document are provided "as is" and without any warranties or guarantees. Reference herein to any specific commercial entity, product, process, or service by trade name, trademark, manufacturer, or otherwise does not constitute or imply its endorsement, recommendation, or favoring by the United States Government, and this guidance shall not be used for advertising or product endorsement purposes.

Purpose

This document was developed in furtherance of the NSA's cybersecurity mission, including its responsibilities to identify and disseminate cyber threats to National Security Systems, Department of Defense, and the Defense Industrial Base, and to develop and issue cybersecurity specifications and mitigations. This information may be shared broadly to reach all appropriate stakeholders.

Contact

Cybersecurity Report Feedback: CybersecurityReports@nsa.gov

General Cybersecurity Inquiries or Customer Requests: Cybersecurity_Requests@nsa.gov

Defense Industrial Base Inquiries and Cybersecurity Services: DIB_Defense@cyber.nsa.gov

Media Inquiries / Press Desk: NSA Media Relations: 443-634-0721, MediaRelations@nsa.gov