Trusted Platform Module (TPM) Use Cases

Executive summary

Trusted Platform Modules (TPMs) are components available on modern computing systems and intended to facilitate several cryptographic, protected storage, and integrity capabilities. NSA recommends acquiring TPMs of version 2.0 or later on devices that support them to be able to leverage their security capabilities for current and future use cases.

DoD Instruction 8500.1 requires inclusion of a TPM for DoD devices subject to DISA STIGs when user credentials and data-at-rest require protection. [1] NSA advocates for several TPM use cases in addition to those required by STIGs, such as for asset management, hardware supply chain security, and boot integrity measurement. Future use cases for the TPM include software supply chain auditing, runtime integrity measurement, and authentication and provisioning to support Zero Trust efforts. TPMs should transition to quantum-resistant cryptography to provide the proper capabilities and assurance for these use cases into the future. As TPM-supporting technologies mature and dependencies are satisfied, these recommended and future use cases may become DoD requirements.

Introduction

Department of Defense (DoD) Instruction 8500.01 directs DoD components to procure computing devices equipped with a Trusted Platform Module (TPM) version 1.2 or later. [1], [2] The directive, originally published in March of 2014, anticipates that TPM capabilities would provide value for computing device identification, authentication, encryption, and integrity measurement within the DoD enterprise infrastructure.

TPMs are explicitly mandated for many devices across DoD as required by DISA's Security Technical Implementation Guides (STIGs). Further, many procurement departments have expanded the adoption of TPMs to include devices even when the current STIGs do not require them. Operating System (OS) vendor compliance requirements for TPMs have also advanced widespread deployment of TPMs. [3] Today, version 2.0 TPMs are commonly found on desktops, laptops, tablets, servers,

and other user devices, not only throughout DoD, but also across most enterprise networks.

DoDI 8500.01 tasks NSA to identify use cases, implementation standards, and plans for DoD integration of TPMs. This document identifies use cases available today and use cases that rely upon unmet dependencies currently in development by the computing industry. NSA-recommended use cases are designed to complement or further strengthen the security impacts of STIG implementation.

These TPM use cases may be affected by requirements relating to quantum-resistant cryptography. For National Security Systems, the requirements and guidance for transitioning to quantum-resistant cryptography are in the latest Committee on National Security Systems Policy 15 (CNSSP 15) and the NSA announcement of the Commercial National Security Algorithms Suite 2.0 (CNSA 2.0) and accompanying frequently asked questions. [4], [5], [6] For other systems, CNSA 2.0 is recommended and for relevant requirements consult the National Institute of Standards and Technology. [7]

TPM form factors

TPMs are available in two physical form factors (sometimes abbreviated to pTPM):

- The first is known as a Discrete or Dedicated TPM (dTPM). Each dTPM is a standalone integrated circuit built to Trusted Computing Group (TCG) specifications and physically attached to the device's mainboard. Most dTPM implementations connect through the Low Pin Count (LPC) bus, Serial Peripheral Interface (SPI) bus, or Inter-Integrated Circuit (I2C) bus. Some may use a daughterboard and pin header to connect to the mainboard.
- The second physical form factor is the Firmware or Integrated TPM (fTPM or iTPM). Each iTPM is built into a host circuit such as the Central Processing Unit (CPU) or Trusted Execution Environment (TEE). Some implementations may place TPM functionality within a Hardware Security Module (HSM) or security coprocessor. Both physical form factors satisfy DoDI 8500.01.

  Note: This document will use and prefer the term iTPM. Treat iTPM and fTPM as synonymous. The term fTPM may cause confusion and is likely to be phased out over time. Confusion often originates from dTPMs having upgradeable firmware.

TPMs are also available in a third form factor — a software form factor, known as a Virtual TPM (vTPM). A Virtual TPM may also be referred to as an Emulated TPM or Software TPM. Each vTPM is a purely software implementation of TCG specifications and may lack any sort of connection to a physical TPM. Pure software TPMs do not satisfy the requirements or intent of DoDI 8500.01.

Cloud infrastructures present a challenge when interpreting DoDI 8500.01's applicability. Some cloud vTPM implementations pass through or otherwise leverage the host hardware's dTPM or iTPM – an implementation that is satisfactory. However, some hypervisors and cloud infrastructures utilize purely software vTPMs that have no link or requirement to use a physical TPM. These software-only solutions may leverage proprietary platform security technologies such as secure enclaves, security coprocessors, and trusted execution environments. Pure software vTPM implementations that have no link to any sort of physical TPM or hardware security technology do not satisfy DoDI 8500.01.

Use cases required by STIGs

User credential protection

TPMs can protect credentials by binding them to the specific keys protected by the TPM of a device such that the credentials are not available for use without the TPM first unwrapping them. DISA STIGs require the use of TPMs when devices are connected to a centralized account management system. For example, devices joined to an Active Directory domain must implement certain TPM-based protections for credentials. The credential protections are provided through Credential Guard on Windows and the Kernel Key Retention Service (keyctl) on Red Hat. Both services must be configured to use the TPM as a root of trust for storage. Certificates and keys must be protected via a TPM storage key and may optionally be bound to a specific TPM integrity state.

Data-at-rest protection

TPMs are widely used to harden Full Disk Encryption (FDE) implementations that encrypt data-at-rest. DISA STIGs require data-at-rest protection for many devices and can be satisfied by using FDE with TPMs. Each TPM employs a unique storage key to protect the FDE key. The FDE key may be further protected with one or more TPM Platform Configuration Registers (PCRs) that bind the key to a specific platform integrity

state. Each PCR summarizes integrity measurements in a specific focus area (e.g., firmware binaries, firmware configuration, Secure Boot values).

Microsoft Windows leverages TPMs as part of its BitLocker FDE solution. [8] BitLocker wraps the Volume Management Key (VMK) with a unique storage key bound to each TPM. By default, integrity data from the system defines the integrity state the device must achieve before the TPM will reveal the VMK. The system integrity data comes from the firmware binary image, Option ROM (OROM) binary images, Master Boot Record (MBR), NTFS boot data, boot loader, and BitLocker policies. To be clear: the host platform must match both the TPM identity and the TPM integrity state for the TPM to unlock the storage key in order to reveal the VMK and be able to decrypt the encrypted data on disk.

Note that dTPM implementations may be vulnerable to LPC or SPI bus probing attacks. Probes could expose the VMK during the reveal process. Utilize BitLocker with PIN as a mitigation until bus encryption solutions become widely available. [9]

Red Hat Enterprise Linux (RHEL) leverages TPMs to harden the Linux Unified Key System (LUKS) or dm-crypt FDE solutions. [10] Both can be configured to utilize TPM-protected keys and bind keys to PCR integrity states. Both provide configuration options that allow administrators to scope protection to the needs of the organization or mission.

Additional recommended use cases

Asset management

TPMs may be leveraged to provide a unique identifier for each device. [11], [12] There are multiple options available to system owners and administrators. Below are three common options:

- Option 1: Each TPM has a unique Endorsement Key (EK) that may be used for unique system identification. The EK will remain consistent and unique unless a Revocable EK (REK) is initialized (but an REK requires TPM owner authorization to access and is blocked by default by the OS). Note that an EK is treated like a serial number or asset tag when used in the manner described in this paragraph. The EK is not able to sign messages or perform attestation like an Identity Key (IK). [13]

Microsoft Windows provides the PowerShell command get-TpmEndorsementKeyInfo which lists a thumbprint value — a SHA-256 hash of the EK. The thumbprint can be used as a unique device identifier. Likewise, RHEL provides tpm_getpubek (TPM 1.x) and tpm2_getekcertificate (TPM 2.x) which can be used to acquire the EK public key. Apply a SHA-256 hash to the EK public key to derive the thumbprint value used by Windows.

- Option 2: Some devices can ship with a manufacturer certificate known as the Platform Certificate (PC). [14] The PC may provide another source of unique identity via the public key certificate or the Platform Serial Number field. This serial number must match the device's serial number usually affixed to the physical system components and recorded in the platform's firmware. When available, the PC provides a standardized and uniform method of acquiring the serial number. The PC may also contain information about the TPM EK's value and a signature originating from the TPM vendor (this may be part of the EK certificate if present).
- Option 3: Device administrators may create a TPM Identity Key (IK) to function as a device's unique identifier within the enterprise. [11] The IK can be accompanied by a certificate — called the Identity Certificate (IC) — and be used for signing operations relating to system integrity attestation. Unlike the EK and PC, TPMs do not ship with an IK at the time of this document's publication. The IK should be created as part of the device deployment into the enterprise infrastructure.

  The IK's strength is that it can provide far more than a static serial number or unique identifier. The IK can be used to respond to proof-of-possession challenges and sign integrity information describing a given device. IKs are bound to a singular TPM and cannot be transferred to another device.

Hardware supply chain auditing

Some system vendors generate a Platform Certificate (PC) for specific business class and server devices. The purchasing organization may need to request inclusion of a PC as part of the procurement process, sometime via a special product identifier or contacting a specific ordering representative. PCs may be preloaded onto the devices prior to shipping or supplied as downloadable content after delivery. [14], [15]

PCs link a device's hardware and firmware with its unique public key certificate, the device serial number, and potentially the TPM EK certificate. Device owners may be able to use this information to link a specific device with a unique order record and device factory configuration record to determine if the device received is what the vendor says it sent. Delta PCs amend the factory information over time to create an audit trail for firmware updates and hardware changes. Use of these certificates empowers an organization to measure whether devices that may outwardly appear to be fully functional have been tampered with during production, transportation, or deployment.

Boot integrity change monitoring

Each TPM stores boot measurements. These measurements may cover firmware binaries, firmware configuration, Secure Boot trust store values, system environment variables, OS kernel integrity, and more depending on device capabilities. Boot measurements are limited in scope and should change infrequently at predictable intervals, making them good candidates for change monitoring. Several Endpoint Detection and Response (EDR) products [16], host-scanning security platforms [17], and open source projects [18] are available to observe boot measurements. Some products monitor for unexpected changes, while others are more sophisticated and provide trustworthiness decisions based on analysis of device measurements.

Use cases dependent on future features

Supply chain auditing

A TPM creates a summary of many individual measurements of the system and stores it as a hash in one of its Platform Configuration Registers (PCRs). Depending on policy configuration, every binary image, interface, library, or application is computed into a hash measurement prior to execution. Each measurement is appended to an applicable PCR depending on the PCR's scope (e.g., firmware, OROM, kernel libraries) and then hashed again as part of the PCR extension process. Understanding the integrity value of a PCR requires auditors to assign trustworthiness to the individual measurements that contributed to the PCR's value. However, few vendors publish measurement hashes that can be used during PCR evaluation.

There are two building blocks that focus on different areas of the supply chain ecosystem that seek to provide a vehicle for vendors to publish known-good measurements. They are:

1. Reference Integrity Manifest (RIM), from the TCG, which is designed to catalog known-good measurement hashes that contribute to TPM PCRs. [19], [20] RIM is particularly relevant to system vendors and OS vendors whose products are most likely collected into PCRs as part of the device boot process. Any number of RIMs from a variety of vendors can apply to a single device depending on the hardware, firmware, and software loaded. RIM is designed to evolve as components are updated and measurements change. RIMs must be signed by the originators of integrity information.

2. Software Bill of Materials (SBOM), advanced by CISA and the National Telecommunications and Information Administration (NTIA), which is designed to catalog measurements of software and related dependencies within the OS user environment. [21], [22] SBOM does not require use of a TPM. SBOM checksum hashes are not guaranteed to directly correlate with TPM PCR measurements of software depending on the configuration of runtime measurement services like Device Health Attestation (DHA) or Integrity Measurement Architecture (IMA). However, the checksum hashes may be used by an SBOM scanning agent (unrelated to TPM) that performs software integrity measurements. The scanner agent may be periodically measured and recorded to a TPM PCR, establishing a trusted integrity chain. SBOMs used as known-good references should be signed by the source responsible for a given piece of software.

In other words, RIM and SBOM allow administrators to compare observed measurements from runtime with provided measurements that the solution vendor(s) indicated are trustworthy. The absence of RIM and SBOM places the onus for determining measurement trustworthiness on administrators. Today, few hardware and software vendors provide RIM and SBOM records to customers.

The unmet dependency that delays leveraging TPM capabilities for supply chain auditing is having a mechanism, presumably managed by system and OS vendors, that automatically calculates and publishes RIM records to device owners in real time or at each product revision. Device owners should begin by requesting that system and OS

vendors publish RIM records regularly as products change. Software purchase and support contracts should request software vendors implement SBOMs to augment or replace package download hashes commonly used today. Software vendors should also implement automation to publish and update SBOM records to match software deployed in production.

Continuous integrity monitoring

Runtime integrity monitoring requires an infrastructure capable of issuing local/network identities to TPM-equipped clients and evaluating their integrity data to make access control decisions. Such an infrastructure must work with TPMs to create identity keys and identity certificates, produce TPM attestation quotes, formulate integrity reports, and utilize RIMs and SBOMs to inform trustworthiness decisions — as part of Network Access Control (NAC) — in real time.

Several products and services offer support for Trusted Network Connect (TNC), PCR auditing, and TPM identity creation. However, a productized system that combines these capabilities with trustworthiness data does not exist at the time of this document's publication. The closest proof-of-concept implementation takes the form of NSA's Host Integrity at Runtime and Startup (HIRS) open source project. [18] HIRS is capable of combining device identity information, attestation quotes, PCRs, and measurements into integrity reports to be evaluated by a remote appraiser. As of this document's publication date, HIRS is not capable of making autonomous access control decisions or automatically integrating known-good measurements provided via RIMs and SBOMs. HIRS can inform device owners of changes to hardware and firmware below the OS, as well as identify changes to privileged software components.

Zero touch provisioning

Zero Touch Provisioning (ZTP) defines an automated mechanism to onboard new devices into an infrastructure. ZTP requires the presence of an Initial Device Identity Key (IDevID) as defined by the TCG and IEEE — use of the EK is not sufficient since it cannot sign challenge responses. [23], [11] The goal of ZTP is to automate the deployment of new trusted devices into infrastructure. ZTP is an enabler for efficiently implementing Zero Trust concepts for new devices within an infrastructure.

A device identity key, Platform Certificate, RIM, SBOM, TPM attestation quote, and infrastructure placement data combine such that the new client is automatically (on first boot) authenticated, integrity checked, and provided an OS image intended for where and how it will be used. Devices with improper credentials, untrusted integrity statuses, or unapproved software are automatically blocked from accessing protected network resources. ZTP represents a mature or advanced implementation of Zero Trust capabilities that requires several enterprise functions to coordinate in real time. However, several of these dependencies currently do not exist or do not coordinate well enough for this capability to be utilized on operational networks. [24]

For Zero Trust Provisioning to be feasible, many infrastructures must be reorganized to include trusted, untrusted, and remediation Virtual Local Area Networks (VLANs). Infrastructure devices tasked with making trustworthiness decisions regarding devices must be 1) ready to recognize the identity of a TPM and parse related device information, 2) capable of evaluating TPM attestation data, 3) informed with RIM and SBOM information, and 4) be capable of executing NAC actions — a combination of capabilities not yet commercially available. Any untrusted device must be evaluated for placement in the trusted or remediation VLANs. Devices in the remediation VLAN should receive new OS images, software packages, updates, or direct administrator interaction to satisfy alerts and deficiencies identified by the infrastructure evaluator — another capability not yet mature enough to support an enterprise Zero Trust implementation.

Conclusion

DoDI 8500.01's TPM requirement was published in 2014. Since that time, government and industry vendors have required inclusion of TPMs in many devices, multiple new versions of the TPM have been released by the TCG, and an entire ecosystem of supporting products and open source projects have become available. DoDI 8500.01 recognizes the importance of a TPM as a trust anchor for modern computing devices and requires its inclusion into applicable devices.

NSA recommends DoD components integrate TPM into their infrastructure for use cases achievable today to further secure DoD missions. Expect more TPM use cases to become requirements in the future as dependencies for more complex use cases are satisfied.

Frequently asked questions

Q: "Is purchasing devices with TPMs sufficient to meet DoD 8500.01? My organization has not implemented a use case."

A: DoD only requires organizations to purchase devices that include TPMs as required by applicable STIGs. NSA and leading OS vendors strongly recommend purchasing and using devices with a TPM. DoD organizations may benefit from TPM use cases even when not required.

Q: "Why are additional use cases 'recommendations' instead of 'requirements'?"

A: NSA believes some use cases may unacceptably impact device availability for some missions as of publication time. Use cases will become required for all DoD components as the TPM ecosystem continues to mature. NSA is working with its partners to create proof-of-concept implementations for some use cases (e.g., HIRS) and encourage vendor capabilities necessary to satisfy dependencies for additional use cases.

Q: "Why do I need a TPM if my organization uses UEFI Secure Boot?"

A: TPM and UEFI Secure Boot are complementary technologies with different advantages and disadvantages. A TPM begins collecting measurements provided to it at the earliest phases of device boot. These measurements include Secure Boot code, certificates, hashes, and other policy artifacts. A TPM also collects measurements describing system components outside the purview of Secure Boot. However, a TPM does not have an active enforcement role — it is an observer by design. An external/third-party verifier, boot loader, kernel, or software agent must query the TPM and then make a decision or take an action based on the TPM's integrity data.

Secure Boot is initialized during the second phase (Pre-EFI Initialization (PEI) phase) of firmware boot. Some Secure Boot implementations and mode configurations check the signature or hash of every boot-time driver, module, and executable while other configurations skip checks to prioritize boot speed. Secure Boot does have an enforcement mechanism that allows it to prevent execution of untrusted binaries. Trust is determined by an ecosystem managed by Microsoft, other vendors, and/or the device's owner. However, Secure Boot lacks the external/third-party verifier capability enabled by TPM attestation.

A TPM and Secure Boot can be combined to cover each other's weaknesses. TPM measurements can provide the integrity state of the firmware that initialized Secure Boot and the values used to define its policy. Secure Boot can actively permit or deny specific binaries the ability to execute at boot. At publication time, trustworthiness data (new trusted signatures and updated revocation records) for Secure Boot is widely available and rapidly updated while TPM lacks equivalent data (which would be provided via RIMs and SBOMs).

Q: "My organization is preparing to make a significant investment in devices. Does NSA have guidance indicating what device capabilities to purchase?"

A: Yes. NSA has published Procurement and Acceptance Testing Guidance [25] designed to help organizations acquire devices capable of implementing all use cases identified in this publication.

Q: "Is there an alternative to a TPM?"

A: If a dTPM or iTPM is not available on a given device and alternatives also do not offer dTPM or iTPM then look for technologies endorsed by the TCG. TPM-like alternatives include the DICE Protection Environment (DPE) and Measurement and Attestation RootS (MARS) initiatives.

Q: Which is better: dTPM or iTPM?

A: Each dTPM is a physical microchip placed on the device's mainboard or attached to it via an add-in card. Common manufacturers of dTPM include Nuvoton, STMicroelectronics, and Infineon. [26] An iTPM is integrated onto a larger package such as the Central Processing Unit (CPU). Manufacturers of iTPM include AMD and Intel. There is no one answer as to which is better. Generally, dTPMs are more exposed to physical eavesdropping and probing techniques affecting their communication channels, but also often have stronger degrees of FIPS certification and tamper protections. iTPMs are far less exposed to eavesdropping at the cost of potentially lower FIPS ratings and exposure to side channel vulnerabilities. Mission needs should dictate which form factor is chosen.

Q: "My organization implements data-at-rest without leveraging a TPM. Is this allowed?"

A: Yes. There are multiple ways to achieve the data-at-rest requirement. Self-Encrypting Drives (SEDs) are a common alternative to TPM-supported FDE. Different missions and organizations have needs that may alter the way STIGs are implemented.

Works cited

[1] Takai, Teresa. March 2014. Department of Defense Instruction Number 8500.01. https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/850001_2014.pdf

[2] Numerous contributors. November 2022. Windows 10 Deployment Scenarios and Tools. https://learn.microsoft.com/en-us/windows/deployment/windows-deployment-scenarios-and-tools

[3] Trusted Computing Group. August 2023. TPM Certified Products. https://trustedcomputinggroup.org/membership/certification/tpm-certified-products/

[4] Committee on National Security Systems. October 2016. CNSS Policy 15. https://www.cnss.gov/CNSS/issuances/Policies.cfm

[5] National Security Agency. September 2022. Announcing the Commercial National Security Algorithm Suite 2.0. https://media.defense.gov/2022/Sep/07/2003071834/-1/-1/0/CSA_CNSA_2.0_ALGORITHMS_.PDF

[6] National Security Agency. April 2024. The Commercial National Security Algorithm Suite 2.0 and Quantum Computing FAQ. https://media.defense.gov/2022/Sep/07/2003071836/-1/-1/0/CSI_CNSA_2.0_FAQ_.PDF

[7] National Institute of Standards and Technology. September 2024. Post-Quantum Cryptography. https://csrc.nist.gov/projects/post-quantum-cryptography

[8] Matarazzo, Paolo & Pamnani, Vinay. June 2023. BitLocker Group Policy Settings. https://learn.microsoft.com/en-us/windows/security/operating-system-security/data-protection/bitlocker/bitlocker-group-policy-settings

[9] Matarazzo, Paolo. November 2023. BitLocker Countermeasures. https://learn.microsoft.com/en-us/windows/security/operating-system-security/data-protection/bitlocker/countermeasures

[10] Peters, Michael. May 2021. What Can You Do with a TPM? https://next.redhat.com/2021/05/13/what-can-you-do-with-a-tpm/

[11] Trusted Computing Group Infrastructure Workgroup. October 2021. TPM 2.0 Keys for Device Identity and Attestation. https://trustedcomputinggroup.org/resource/tpm-2-0-keys-for-device-identity-and-attestation/

[12] Trusted Computing Group Infrastructure Workgroup. August 2015. TPM Keys for Platform Identity for TPM 1.2. https://trustedcomputinggroup.org/resource/tpm-keys-for-platform-identity-for-tpm-1-2-2/

[13] Trusted Computing Group. July 2024. TCG Platform Certificate Profile. https://trustedcomputinggroup.org/resource/tcg-platform-certificate-profile/

[14] Prowess Consulting. December 2022. What is the Best Supply-Chain Solution to Verify Server Configurations? https://www.delltechnologies.com/asset/ko-kr/products/cross-company/industry-market/prowess-dell-supply-chain-solutions-vs-hpe-whitepaper.pdf

[15] HP Wolf Security. June 2023. HP Platform Certificate Customer Frequently Asked Questions (FAQ). https://h20195.www2.hp.com/v2/GetDocument.aspx?docname=4AA8-3108ENW

[16] Ionescu, Alex. May 2019. CrowdStrike Falcon: First Endpoint Protection to Integrate Firmware Attack Detection Capability. https://www.crowdstrike.com/blog/crowdstrike-first-to-deliver-bios-visibility/

[17] Loucaides, John. June 2020. Ensuring Device Security in Federal Environments. https://eclypsium.com/solutions/firmware-security-for-enterprises/

[18] National Security Agency. February 2022. HIRS (Host Integrity at Runtime and Startup). https://github.com/nsacyber/HIRS/

[19] Trusted Computing Group. November 2020. TCG Reference Integrity Manifest (RIM) Information Model. https://trustedcomputinggroup.org/resource/tcg-reference-integrity-manifest-rim-information-model/

[20] Trusted Computing Group PC Client. April 2024. TCG PC Client Reference Integrity Manifest Specification. https://trustedcomputinggroup.org/resource/tcg-pc-client-reference-integrity-manifest-specification/

[21] Cybersecurity and Infrastructure Security Agency. November 2023. Software Bill of Materials (SBOM). https://www.cisa.gov/sbom/

[22] National Security Agency. January 2024. Recommendations for Software Bill of Materials (SBOM) Management. https://media.defense.gov/2023/Dec/14/2003359097/-1/-1/0/CSI-SCRM-SBOM-MANAGEMENT.PDF

[23] Institute of Electrical and Electronics Engineers. August 2018. 802.1AR: Secure Device Identity. https://1.ieee802.org/security/802-1ar/

[24] National Security Agency. October 2023. NSA Shares Recommendations to Advance Device Security Within a Zero Trust Framework. https://www.nsa.gov/Press-Room/Press-Releases-Statements/Press-Release-View/Article/3562460/nsa-shares-recommendations-to-advance-device-security-within-a-zero-trust-frame/

[25] National Security Agency. September 2023. Procurement and Acceptance Testing Guide for Servers, Laptops, and Desktop Computers. https://media.defense.gov/2023/Sep/28/2003310132/-1/-1/0/CSI_PROCUREMENT_ACCEPTANCE_TESTING_GUIDE.PDF

[26] Trusted Computing Group. August 2024. TPM Certificate Products. https://trustedcomputinggroup.org/membership/certification/tpm-certified-products/

Disclaimer of endorsement

Purpose

This document was developed in furtherance of NSA's cybersecurity missions, including its responsibilities to identify and disseminate threats to National Security Systems, Department of Defense, and Defense Industrial Base information systems, and to develop and issue cybersecurity specifications and mitigations.

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