Advancing Zero Trust Maturity Throughout the Visibility and Analytics Pillar

Executive summary

In the ever-expanding landscape of cybersecurity, threats manifest in various forms and often infiltrate systems discreetly. The constant risk of intrusion underscores the critical importance of swift detection and mitigation.

This cybersecurity information sheet (CSI) centers on the visibility and analytics aspect of the Zero Trust (ZT) model, emphasizing the significance of comprehensively observing data characteristics and events within an enterprise-wide environment. Prioritizing cyber-related data analysis aids in informing policy decisions, facilitating response actions, and constructing a risk profile to proactively fortify security measures.

Visibility and analytics form the cornerstone of any ZT strategy, empowering organizations to harness infrastructure, tools, data, and techniques for proactively mitigating risks and for rapid identification, detection, and response to emerging cyber threats. Evolving from traditional signature-based approaches, detection (visibility and analytics) and response capabilities are increasingly adopting behavior-based methodologies to combat the sophistication of modern cyber threats. This pillar highlights the benefits of continuous monitoring and provides insights essential for identifying and mitigating potential security risks to assure that only authorized users and devices access sensitive resources.

This CSI offers recommendations for advancing visibility and analytics within the ZT framework. It explains how these capabilities seamlessly integrate into a comprehensive ZT framework as detailed in the NSA publication, Embracing a Zero Trust Security Model. [1] National Security System (NSS), Department of Defense (DoD), and Defense Industrial Base (DIB) stakeholders can leverage this guidance in conjunction with complementary resources to enhance visibility and analytics through the implementation of outlined capabilities.

Introduction

The advent of generative artificial intelligence, capable of mimicking human communication patterns with remarkable speed, has exacerbated the phishing threat. [2] To effectively combat phishing, it is crucial to delve into various indicators and understand the requisite visibility needed to capture relevant information comprehensively. Once this data is gathered, the fusion across disparate sources becomes imperative to stitch together the attack pattern triggered by the initial phishing attempt. Traditional IT architectures often falter in detecting this threat due to their inability to adapt and respond swiftly to the dynamic nature of phishing techniques, underscoring the urgent need for enhanced visibility and analytics capabilities.

Researchers have learned repeatedly that the actions of malicious cyber actors are rarely impulsive and unplanned. These entities often act in intentional and deliberate ways; a concept captured in the MITRE® ATT&CK® knowledge base focused on the specific taxonomy of adversarial actions. [3] Detecting and identifying potential threats requires both human and technological elements to understand the entirety of the network, to detect anomalous changes that may signal intrusion, and to react to an incident expediently and properly should one occur.

The individual concepts in this CSI are not entirely novel, but are presented in a manner that identifies their unique and interdependent roles in supporting cybersecurity through a complete Zero Trust (ZT) framework. They work together to uphold the key principles of ZT, while focusing more directly on the visibility and analytics pillar of capabilities and activities as defined in the DoD Zero Trust Strategy. [4]

Audience

This CSI provides guidance primarily intended for National Security System (NSS), Department of Defense (DoD), and Defense Industrial Base (DIB) networks, 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) [5] and the Cybersecurity and Infrastructure Security Agency (CISA). [6] This guidance is compatible with the DoD's Cybersecurity Reference Architecture (CSRA) Version 5.0, [7] the DoD Zero Trust Reference Architecture (ZT RA) Version 2.0, [8] and the DoD Zero Trust Strategy, [4] referenced at the end of this document.

Background

The President's Executive Order on Improving the Nation's Cybersecurity (EO 14028) [9] and National Security Memorandum 8 (NSM-8) [10] direct the Federal Civilian Executive Branch (FCEB) agencies and NSS owners and operators to develop plans to adopt a ZT cybersecurity framework.

In Embracing a Zero Trust Security Model, the concept of ZT is defined and contextualized along with the undergirding principles of the seven pillars. The pillars are composed of capabilities that enable progressive maturity across a comprehensive ZT framework. [1] The capabilities described in this CSI 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.

This CSI offers guidance in accordance with the DoD ZT RA maturity scale of preparation, baseline, intermediate, and advanced. It also complements the ZT Portfolio Management Office (PfMO) guidance of "target" and "advanced" levels. [4]

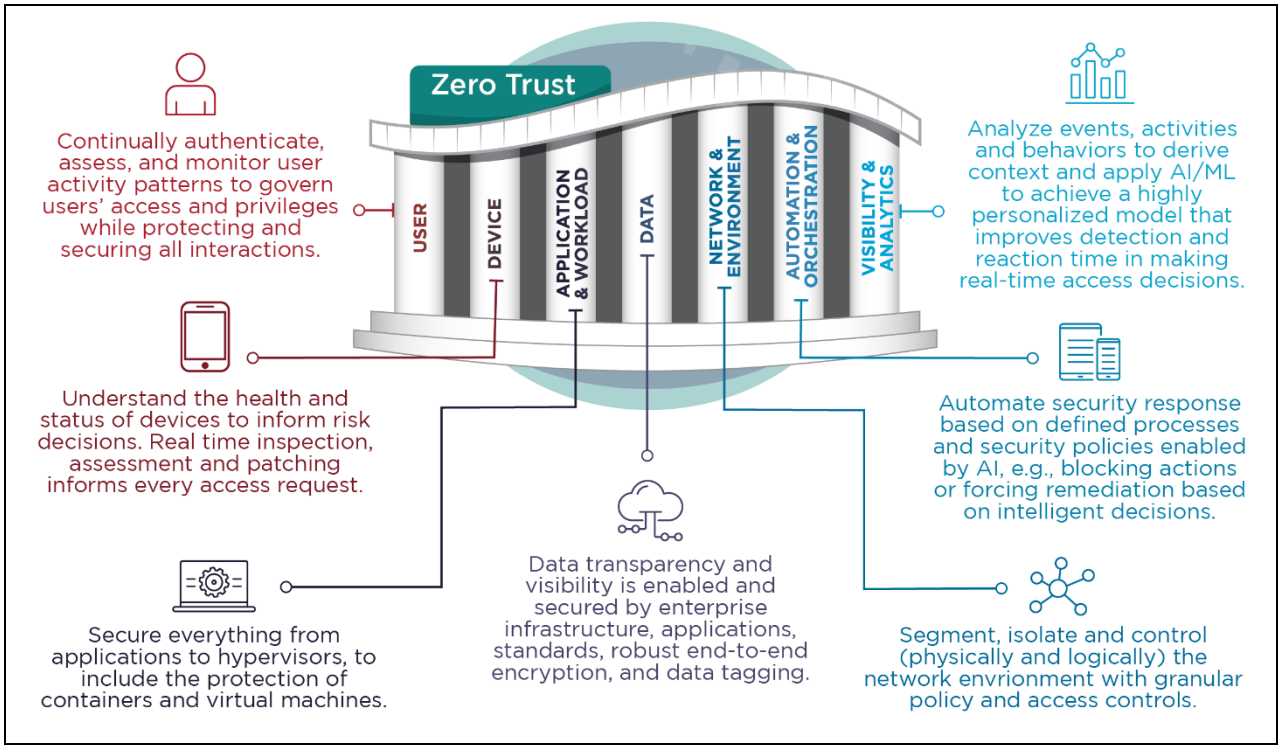

Figure 1: Description of the seven pillars of Zero Trust

Figure 1 displays the seven ZT pillars, including the visibility and analytics pillar. The capabilities and milestones for the visibility and analytics pillar of the ZT maturity model will be described in detail throughout this document. The seven ZT pillars are not independent as each pillar depends on or aligns with the capabilities in the other pillars.

Visibility and analytics pillar

According to the DoD ZT RA, visibility improves detection of anomalous behavior and provides the ability to make dynamic changes to security policy and real-time access decisions. [8] The Cybersecurity and Infrastructure Security Agency (CISA) Zero Trust Maturity Model Version 2.0, which aligns with NSS guidance, refers to visibility as observable artifacts that result from the characteristics of and events within enterprise-wide environments. [6] The visibility and analytics pillar in the ZT framework provides health, status, performance, behavioral, and threat insights across the infrastructure by observing real-time communications and security-relevant activities happening across all network components. This pillar is necessary for baselining data needed for the other pillars because it leverages log data from all of them to build a comprehensive picture of what is happening across the network. Not only does this pillar depend on the other pillars, but then it employs its own capabilities to bring together and make sense of the collected data, producing critical insights into anomalous and potentially malicious activity.

This is accomplished through the following key capabilities to be discussed in greater detail below:

- Logging all relevant activity
- Centralized security information and event management
- Security and risk analytics
- User and entity behavior analytics
- Threat intelligence integration
- Automated dynamic policies

These capabilities ensure that proper implementation of the visibility and analytics pillar will facilitate analysis of events, activities, and behaviors and feed artificial intelligence and machine learning technology for additional actionable analytics. Success in the visibility and analytics pillar ensures that detection and response to cybersecurity threats is prompt, and allows for threat hunting, forensics, accelerated investigations, and support to compliance initiatives.

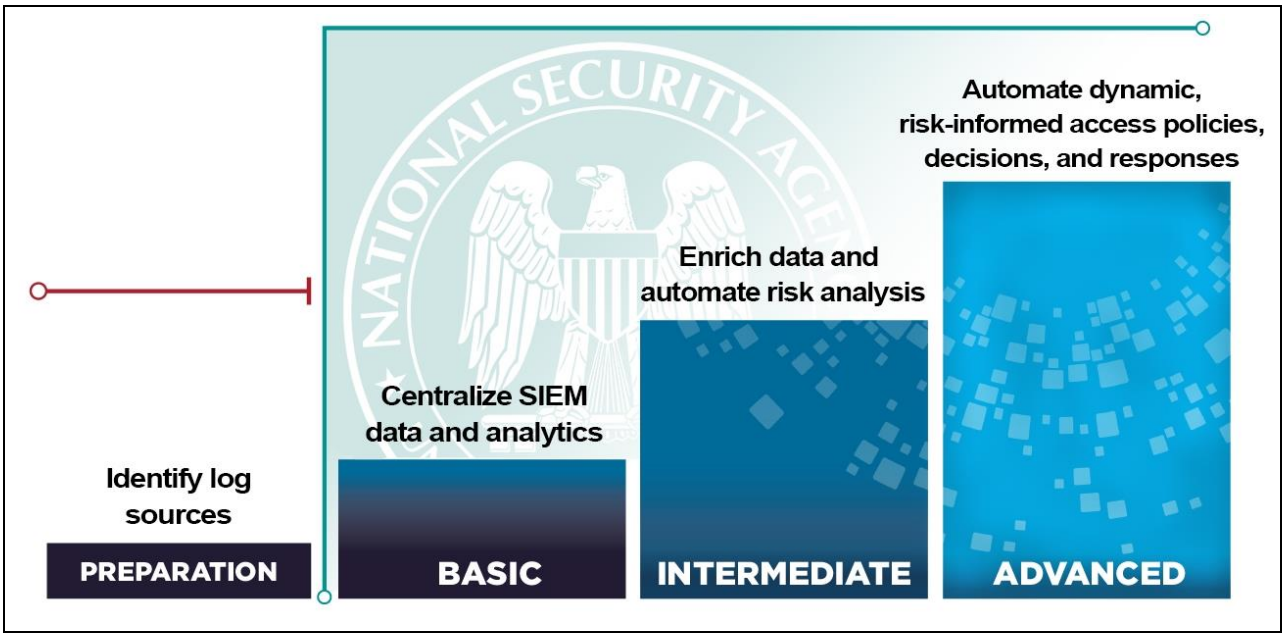

Figure 2: Zero Trust visibility and analytics pillar maturity

Logging all relevant activity

The first step in making sense of what is happening on the network is to capture relevant activity logs from network devices (e.g., switches, routers, wireless access points, network monitoring software), user devices and applications (e.g., office products, development tools, operating systems), authentication services (e.g., identities and credentials used, account directories, federation gateways, etc.), and resources and data access (e.g., file systems, hosted services, email services, websites, etc.). Given the amount of log data this could entail, capture only pertinent activities occurring across the environment with priority given to the right viewpoints rather than indiscriminately collecting everything. Begin by identifying all available log sources and potential log events, and then choosing the ones to capture that have relevant, useful information. These logs should be gathered into a central repository, like a SIEM, and stitched together into a workflow that can be analyzed for indicators of suspicious or malicious activity.

Traditionally, this data has been used in isolated computer security or information technology initiatives (e.g., operations, audits, user behavior, regulatory compliance) as snapshot activities, such as annual assessments or penetration testing, or as supplementary information sources on an as-needed basis. Collecting and understanding activities via logs allows for the baselining and determination of what is

considered normal activity on the network. Once established, any deviation from the norm can be investigated to determine if an activity is benign (such as an accidental misconfiguration) or malicious. The logs must be collected, stored, standardized per industry standards, and analyzed continuously for network security purposes to always have a view of what is occurring on the network.

Table 1: Logging all relevant activity maturity

Preparation	Basic	Intermediate	Advanced
Organizations collect and store all pertinent network, data, application, device, and user logs. A network baseline is established, and an update schedule to ensure snapshots are kept current.	Logs are parsed, normalized, and shared with the appropriate network defenders and security operations center (SOC).	Logs relevant to hunt and incident response operations are identified and prioritized for retention and analysis. Additional logging is processed as needed.	Logging activities are automated to automatically ingest, format, and normalize relevant logs from new sources. This provides a more complete picture of network activities and user behavior and facilitates security protections automation at scale.

Centralized security information and event management

Organizations should utilize security information and event management (SIEM) technology to aggregate and centralize security information about a network, correlate events from different sensor viewpoints, and perform analysis on the data to generate alerts about potential cybersecurity threats within a network environment. A SIEM manages this data by providing interactive user dashboards with actionable playbooks, fast indexing for search queries, and advanced processing capabilities (e.g., machine learning or high performance computing (HPC) capabilities to manage events at scale) for preprocessing and detection of threats. The SIEM's objective is to identify security threats and enable them to be mitigated before they have an opportunity to cause harm to an organization.

Table 2: Centralized security information and event management (SIEM) maturity

Preparation	Basic	Intermediate	Advanced
Network defenders and security operations centers (SOC) monitor, detect, and analyze data logged into a SIEM tool manually.	User and device baselines are created using security controls and integrated with the SIEM. Analytics and alerting within the SIEM is matured to support more advanced data points (e.g., enrichment with cyber threat intel, automated baseline comparisons, etc.)	More data is ingested, processed, and analyzed automatically for anomalous user behavior, malicious activity alerting, and automating collecting related information relevant for incident response to common threat events.	Use of machine learning to help identify unknown threat behaviors. Use of high performance computing (HPC) capabilities for processing data faster and at scales not previously achievable.

Security and risk analytics

Organizations aim to develop analytics that assess risk to their data and infrastructure using information about their configurations, network assets, and known exploited vulnerabilities. The risk picture should leverage information about security capabilities implemented to protect systems and data as factors that may reduce the initial risk assessment. The risk estimation should also factor in the value and criticality of assets in the network, such as important files, critical servers, essential services, or highly privileged users, when determining the potential impact of possible incidents. External data, such as Common vulnerabilities and exposures (CVEs), known exploits (ExploitDBs), and the Common Vulnerability Scoring System (CVSS), should be utilized to enrich dynamic risk scores as the threat environment changes.

Table 3: Security risk analytics maturity

Preparation	Basic	Intermediate	Advanced
Data, such as vulnerabilities from CVEs,	Known vulnerabilities, configuration	Vulnerability information from CVEs, known exploited vulnerability	A comprehensive risk assessment scoring system is

Preparation	Basic	Intermediate	Advanced
exploits from ExploitDB, etc., are collected and used to develop CVSS scores.	information, and asset data are combined to develop a simple risk assessment.	data, asset information, and risk mitigation measures are analyzed to develop locally adjusted CVSS scores and more comprehensive risk assessments. Initial insider threat risks are evaluated.	applied to networks, including information about the sensitivity of data and criticality of assets and services. Additional information regarding new vulnerabilities is potentially scraped from authoritative sources to better determine risk. An automated risk assessment system is established.

User and entity behavior analytics

User and entity behavior analytics (UEBA) utilize log data to detect abnormal behaviors occurring on the network. UEBA usually attempts to correlate across many log sources to find unusual behaviors, especially since single or limited log sources are often insufficient to detect malicious behaviors by cyber actors who try to hide within the noise of normal network activities and fluctuations. Behavior analytics were first described in the user pillar with analysis of basic user patterns. UEBA within organizations completes its expansion here by taking collected data and using traditional and machine learning (ML) artificial intelligence (AI) approaches to analyze the enormous set of network activities and pick out aberrations that may indicate malicious activity. Initially, AI-based detections are supervised, but using advanced techniques, such as neural networks, UEBA operators may not need to be part of the learning process. These models, however, will always be probabilistic; models that cannot be fully explained will need to be backed up with evidence, and human operators will always be necessary for operational decisions.

The objective of UEBA is to identify abnormal and potentially dangerous behavior indicative of threats. Identification of insider threats is the key objective in locating malicious insiders or hackers who manipulate compromised insider credentials. These insider threats have the capability to escape notice by other security tools since they act with legitimate permissions or mimic authorized activity on a network.

MITRE's D3FEND™ taxonomy of cybersecurity countermeasures includes an entire technique category of User Behavior Analysis [D3-UBA] with a dozen sub-techniques for different methods of identifying anomalous behaviors that could indicate potential malicious activity. [11]

Table 4: User and entity behavior analytics maturity

Preparation	Basic	Intermediate	Advanced
UEBA baselining and profiling strategy is designed. High-risk areas are identified, and specific use-cases are established.	Organizations initially employ analytics to profile and baseline activity and behaviors of users and entities and to correlate those baselines to inform anomaly detection.	Abnormal and potentially dangerous user and device behavior threats are automatically identified.	Advanced automated AI-based analytics supporting detection of anomalous users, devices, and other entity actions and advanced threats are created and used.

Threat intelligence integration

Threat intelligence helps organizations enrich their awareness and develop tools and techniques aimed at threats or activities by threat actors. Threat intelligence information includes threat indicators; tactics, techniques, and procedures (TTPs); security alerts; threat intelligence reports; tool configurations, and other sources. Threat intelligence helps organizations prioritize security events and alerts based on severity and relevance, identify potential threats that may have gone previously undetected, enhance AI/ML model accuracy, and inform decision-making processes for security analysts. Consider using indicator sharing services from government or sector sharing initiatives, as well cybersecurity advisories (CSAs) published by NSA or alerts by other cybersecurity authorities. Many advisories and alerts include indicators that can be used by cybersecurity tools and SIEMs. Additionally, threat information informs network

owners of the latest adversary TTPs and can be aggregated, transformed, analyzed, or interpreted to help set priorities, including for logging and mitigations.

Table 5: Threat intelligence integration maturity

Preparation	Basic	Intermediate	Advanced
The organization defines its threat intelligence requirements, incorporating critical assets, known threats, use of threat intelligence data, etc.	Threat Intelligence Platforms (TIPs) are implemented, data is normalized for standardization and machine-readability. SOCs gather threat intelligence information and streams.	SOCs combine threat intelligence information to focus on identities, characteristics, and TTPs that can be matched by data collected in the SIEM.	Resultant threat intelligence data is directly integrated into other SIEM data to enhance automated detection and response efforts.

Automated dynamic policies

Informed by the other capabilities in the visibility and analytics pillar, organizations establish dynamic policies by utilizing previous rule-based access to teach AI/ML algorithms to make automated access decisions to various resources. Ongoing monitoring of security posture, risk assessments and scoring, and automated patch management are all used in this process.

One of the ultimate goals of ZT is to establish dynamic policy automation modeled using real-time security profiles and have those policies continuously adapt themselves based on the changing threat environment and confidence scores. The same applies to other policies beyond access control as well, such as patching and configuration policies. Depending on the organization's policies and risk tolerance, dynamic policies may require people to approve policy changes, often through pre-approved policy options that can be implemented dynamically.

Table 6: Automated dynamic policies maturity

Preparation	Basic	Intermediate	Advanced
Existing policies are reviewed and determinations made regarding which policies are suitable for automation. New dynamic policies developed as needed to address emerging threats.	AI/ML solutions are established and applicable policies are automated on limited datasets determined by risk level.	Organization AI/ML solutions dynamically and automatically update security profiles and device configuration through continuous security posture monitoring, risk and confidence scoring, and automated patch management.	User access established based on automated, real-time security profiles where feasible and appropriate. Continuous policy development based on evolving conditions and risk and confidence scores.

Summary of guidance

Listed below is a summary of relevant guidance specific to the visibility and analytics ZT pillar:

- Develop a plan for log collection and integration of the highest priority logs (e.g., firewalls, Endpoint Detection & Response, Active Directory, switches, routers, etc.) within a common SIEM, followed by lower priority logs. Adopt an open industry-standard log format.
- Develop analytics and detection capabilities within the SIEM.
- Use external threat intelligence and vulnerability information to augment and enrich detection capabilities of risks and potential threat activities.
- Incorporate AI/ML when possible into these analytics processes to improve scale, scope, and efficiency.

The information presented herein is not a definitive guide with a standardized solution that fits all organizational needs, but rather suggestions and considerations for implementing effective visibility and analytics within the ZT framework. Discovering and defining the organization's mission and identifying the supporting assets which need to be secured will help construct a clearer picture of the as-is architecture compared against recommendations in the seven ZT pillar CSIs. This comparison will assist all stakeholders to identify organizational risks and gaps, while forming a mature ZT

architecture for specific organizations. Each organization will need to evaluate ZT security requirements to determine which are applicable to them. The overall government goal is to develop a ZT roadmap with strategies which will align with the organization's ZT goals.

Further guidance

NSA is assisting customers across the NSS community that are piloting ZT capabilities and coordinating ZT activities with the NSS community, NIST, CISA, and DoD. NSA is also assisting in the development of additional ZT guidance to support system developers through the challenges of integrating ZT within NSS, DoD, and DIB environments. Upcoming additional guidance will help organize, guide, and simplify incorporating ZT principles and designs into enterprise networks.

Works Cited

[1] National Security Agency. Embracing a Zero Trust Security Model. 2021. https://media.defense.gov/2021/Feb/25/2002588479/-1/-1/0/CSI_EMBRACING_ZT_SECURITY_MODEL_UOO115131-21.PDF

[2] Z. Amos. Phishing in 2024: Here's What to Expect. 2024. https://cybersecurity-magazine.com/phishing-in-2024-heres-what-to-expect/

[3] MITRE.org. MITRE ATT&CK. 2024. https://attack.mitre.org/

[4] Department of Defense. DoD Zero Trust Strategy. 2022. https://dodcio.defense.gov/Portals/0/Documents/Library/DoD-ZTStrategy.pdf

[5] National Institute of Standards and Technology. NIST Special Publication 800-207. 2020. https://csrc.nist.gov/publications/detail/sp/800-207/final

[6] Cybersecurity and Infrastructure Security Agency. CISA Zero Trust Maturity Model Version 2.0. 2023. https://cisa.gov/zero-trust-maturity-model

[7] Department of Defense. DoD Cybersecurity Reference Architecture (CSRA) Version 5.0. 2023. https://dodcio.defense.gov/Portals/0/Documents/Library/CS-Ref-Architecture.pdf

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

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

[10] The White House. National Security Memorandum 8: Improving the Cybersecurity of National Security, Department of Defense, and Intelligence Community Systems. 2022. 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/

[11] MITRE.org. MITRE D3FEND. 2024. https://d3fend.mitre.org/

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 products, 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.

Trademarks

ATT&CK and MITRE and are registered trademarks of The MITRE Corporation. D3FEND is a trademark of The MITRE Corporation.

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. 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