infoblox.

National Institute of Standards and Technology
United States Department of Commerce
100 Bureau Drive
Gaithersburg, MD 20899
cyberframework@nist.gov

Subject: Infoblox Comments to NIST Cybersecurity Framework 2.0 Public Draft and NIST Cybersecurity Framework 2.0 Core with Implementation Examples Discussion Draft

Thank you for the opportunity to provide our comments and recommendations to the Drafts. We commend NIST for its efforts to update and improve CSF based on community input and changing cybersecurity challenges and best practices.

We provide the following comments guided by NIST's zero trust principles as well as our experience in providing real-time visibility and control over who and what connects to an organization's network by uniting networking and security intelligence.

1. Current and historical attributable metadata should be added to ID.AM and DE.AE Categories

"Zero trust presents a shift from a location-centric model to an identity, context, and data-centric approach with fine-grained security controls between users, systems, applications, data, and assets that change over time."[1] To move from a Traditional starting point to Initial stage under NIST's Zero Trust Maturity (ZTM) journey, NIST recommends an enterprise "collects as much information as possible about the current state of assets, network infrastructure, and communications and uses it to *improve its security posture (emphasis added)*."[2]

It would be helpful for NIST to provide some recommendations for the types of information an organization should prioritize on its ZTM journey. We believe that key attributable metadata, both current and historical, associated with an organization's assets would have the highest impact on rapid incident response and proactive vulnerability management. This applies to not only physical assets, but also to virtual assets and IoTs, wherever located.

To that end, we propose the following changes to the ID.AM and DE.AE categories, each as shown in the underlined red text below:

[1] https://www.cisa.gov/sites/default/files/2023-04/CISA_Zero_Trust_Maturity_Model_Version_2_508c.pdf
[2] https://www.cisa.gov/sites/default/files/2023-04/CISA_Zero_Trust_Maturity_Model_Version_2_508c.pdf

Category	Subcategory	Implementation Examples
Asset Management (ID.AM): Assets (e.g., data, devices, software, systems, facilities, people) that enable the organization to achieve business purposes and current and historical key attributable metadata are identified and managed consistent with their relative importance to organizational objectives and the organization's risk strategy	**ID.AM-05:** Assets are prioritized based on current and historical classification, criticality, resources, and impact on the mission	**Ex3:** Track the asset priorities on a real time basis and update them dynamically ~~periodically~~or when significant changes to the organization occur Ex4: Maintain a history of key attributable metadata allocated to each asset and resource
Adverse Event Analysis (DE.AE): Anomalies, indicators of compromise, and other potentially adverse events are analyzed to characterize the events and detect cybersecurity incidents (formerly DE.AE, DE.DP-02)	**DE.AE-07:** Cyber threat intelligence and other contextual information are integrated into the analysis	**Ex2**: Securely provide information from asset inventories and current and historical key attributable metadata (e.g., network location, IP address, ownership, data classification) to detection technologies, processes, and personnel

The following explains the reasons behind the proposed addition of "current and historical attributable metadata":

A. What constitutes key attributable metadata for an asset?

We refer to key attributable metadata as data that can accurately and quickly improve the security posture of an organization's assets, network infrastructure, and communications in line with NIST's ZTM guidance.

For devices (both physical and virtual), key attributable metadata would include network location, IP address, ownership, and data classification. We believe that key attributable metadata, both current and historical, would have the highest impact on rapid incident response and proactive vulnerability management.

Incident Detection and Response: Consider the following scenario. At 5:00 p.m., various assets of an organization were compromised by malware, including a laptop of an IT administrator, a new IoT device for greenhouse gas monitoring, and a server from a recent acquisition containing crucial client data. Amidst the many alerts that day, the SOC needed to promptly identify and quarantine the compromised assets, prioritizing the laptop and the server due to their classification.

The above-mentioned key attributable metadata (i.e., network location, IP address, ownership and data classification) can help an organization's SOC team to quickly identify who, what and when: which device was compromised, who owned the device, what threat was involved, what activities occurred on the device, what data was stored on or accessed from the device, when the attack reached the compromised device, and last but not least, which device and what remedial actions to prioritize.

Vulnerability Management: Such key attributable metadata is also important in proactive vulnerability remediation. An organization armed with such metadata would be able to quickly identify the specific assets associated with vulnerable IP addresses, reducing risk.

B. Why is current and historical metadata necessary?

Mapping a SOC alert to a device can be difficult and time-consuming because some asset metadata such as IP address, network location and ownership are often dynamic. Multiple assets could be assigned the same network location, IP address or ownership at different times or even at the same time. A cyber alert typically starts with an IP address, and yet the same IP address can be reassigned to a different device and/or owner. Mapping the IP address to the compromised asset, specifically at the time of the attack, requires tracking key attributable metadata on a real-time basis, together with a history of its allocation to each asset and resource.

Continuing the example set forth in **A** above, by 5:01 p.m., the IP address of the privileged laptop was reassigned to another user due to IP reallocation; post-reboot, the server acquired a different IP address. The SOC team had to coordinate with the network services team to research the IP address history, network location, access logs, assigned users, and other necessary metadata. This manual process ended up taking several days with the assistance of multiple employees for each alert. Automating the allocation of real-time key attributable metadata, both current and historical, and making it accessible to both the SOC team and the network service team can expedite research and remediation efforts.

Connecting network intelligence with threat intelligence will provide an organization with one single view of all its assets connected to the network in the most cost-efficient way on a real time basis. Adding "current and historical key attributable metadata" that improves security posture will provide the much needed guidance around what type of information an organization should prioritize to collect on its ZTM journey.

2. Protective DNS and Suspicious Activity should be added to DE and PR Categories

Domain Name Services (DNS) resolves human-readable host names to Internet Protocol (IP) addresses. As a fundamental network service, DNS has to be left open to enable internet connections, and as a result, it has been used by threat actors as a strategic vehicle to send malware and conduct data exfiltration, command and control (C2), etc. Per the Cybersecurity and Infrastructure Security Agency (CISA), "DNS Infrastructure is a common threat vector for attack campaigns."[3] It serves as a vehicle just as common as (or more than) email, web, file sharing and collaboration services.

DNS has been used as an attack vector since its inception. In the December 2020 Solarwinds attack, hacked networks were communicating with domains the attackers set up as a Command and Control (C2). This pattern has repeated itself over the years and continues today, regardless of target and motive. On August 27, 2023, Retool's employees fell victim to phishing messages containing a malicious URL link that mimicked Retool's own internal identity portal. The attacker was able to use the stolen MFA tokens to access several internal systems and take over 27 customer accounts. Today, large quantities of registered domains are being generated almost at once in support of various networks, both for legitimate uses, and for malicious uses to distribute malware through malvertising campaigns and other similar schemes. Threat actors are leveraging not only improved CPU and GPU performance but also generative AI and quantum computing to evade detection.

A. Why do we need to add Protective DNS?

Recognizing the pivotal role of DNS in such attacks, governments around the world, including CISA, National Security Agency (NSA), UK National Cyber Security Center, and Australian Cyber Security Centre, have championed the use of Protective Domain Name Systems (Protective DNS) and made it available to its agencies.[4] In addition, CISA and NSA recommend Protective DNS use by private sectors. They also encourage enterprise networks to select "enterprise PDNS services that provide malicious activity alerts, enterprise dashboard views, historical logging and analysis, and other enterprise-focused features".[5] A PDNS solution should be able to detect DNS tunneling, anomalous behaviors in DNS packet headers and related traffic, signatures associated with exploitation techniques across DNS sessions, in both the outbound query and the inbound response, and most importantly–alert and stop it in the quickest time possible.

Adding "Protective" in front of "DNS service" in the Implementation Ex.3 under the PR.PS Category will align NIST CSF with the guidance from CISA and NSA.

---

[3] https://www.cisa.gov/news-events/news/cisa-launches-its-protective-dns-resolver-general-availability-federal-agencies

[4] https://media.defense.gov/2021/Mar/03/2002593055/-1/-1/0/CSI_Selecting-Protective-DNS_UOO11765221.PDF; https://www.ncsc.gov.uk/information/pdns; https://www.minister.defence.gov.au/media-releases/2021-10-14/new-cyber-guard-government-data

[5] https://media.defense.gov/2021/Mar/03/2002593055/-1/-1/0/CSI_Selecting-Protective-DNS_UOO11765221.PDF

B. Why do we need to add Suspicious Activity?

Furthermore, CISA and NSA specify that "PDNS service may take several actions to respond to a malicious or ***suspicious*** domain name query" (*emphasis added)*.[6] The specific reference of "suspicious" is important because it takes months, sometimes years, for a "suspicious domain" to be confirmed as malicious. Many threat actors strategically age the domains before leveraging them, or utilize dynamic DNS which allows their IP addresses to change rapidly. For example, a malicious domain used in the Solarwinds attack, avsvmcloud[.]com, was registered on July 25, 2018. It was active for 2.5 years, and used by malware for 6 months, before being publicly identified as a malicious domain on December 14, 2020. Infoblox finds that it takes on average approximately 60-90 days before suspicious domains are publicly associated with malware campaigns. It is important for an organization to detect and block these domains before they are actively used in a malicious campaign. When organizations block domains that are "suspicious" rather than known malware, they are protected before the exact nature of the threat is known.

As a matter of fact, NIST CSF 2.0 Draft already reflects the preventative mindset by extending beyond "known malicious" domains or activities. For example, Implementation Ex1 under DE.AE-02 specifically requires the continuing monitoring of "log events for known malicious and ***suspicious*** activity" *(emphasis added).* Implementation Ex2 under PR.AT-01 calls for training users "to recognize social engineering attempts and other common attacks, report attacks and ***suspicious*** activity" *(emphasis added).* There are 8 references to "***potentially*** adverse events" *(emphasis added)* in the DE.CM and DE.AE Categories.

To conclude, we believe NIST's zero trust framework supports the inclusion of a Protective DNS capability resilient to both known and ***suspicious*** threats in the DE and PR Categories, each as shown in the underlined red text below. These changes are consistent with the "low regret" methodology published by Johns Hopkins APL,[7] CISA and NSA's guidance mentioned above,[8] and the emphasis on detecting suspicious activity in other sections of the NIST CSF 2.0 Draft.

Category	Subcategory	Implementation Examples
Continuous Monitoring (DE.CM): Assets are monitored to find anomalies, indicators of compromise, and other potentially adverse events	**DE.CM-09:** Computing hardware and software, runtime environments, and their data are monitored to find potentially adverse events (formerly PR.DS- 06,	**Ex1:** Monitor email, web, file sharing, collaboration services, DNS, and other common attack vectors to detect malware, phishing, data leaks

[6] https://media.defense.gov/2021/Mar/03/2002593055/-1/-1/0/CSI_Selecting-Protective-DNS_UOO11765221.PDF
[7] https://github.com/JHUAPL/Low-Regret-Methodology/blob/main/README.md#introduction
[8] https://media.defense.gov/2021/Mar/03/2002593055/-1/-1/0/CSI_Selecting-Protective-DNS_UOO11765221.PDF

	PR.DS-08, DE.CM-04, DE.CM-05, DE.CM-07)	and exfiltration, and other potentially adverse events
Platform Security (PR.PS): The hardware, software (e.g., firmware, operating systems, applications), and services of physical and virtual platforms are managed consistent with the organization's risk strategy to protect their confidentiality, integrity, and availability	**PR.PS-05:** Installation and execution of unauthorized software are prevented	**Ex3:** Configure platforms to use only approved Protective DNS services that block access to known malicious and suspicious domains and IP addresses
Technology Infrastructure Resilience (PR.IR): Security architectures are managed with the organization's risk strategy to protect asset confidentiality, integrity, and availability, and organizational resilience	**PR.IR-01:** Networks and environments are protected from unauthorized logical access and usage (formerly PR.AC-03, PR.AC-05, PR.DS-07, PR.PT- 04)	**Ex2**: Logically segment organization networks from external networks, and permit only necessary communications to ~~enter~~ transit the organization's networks ~~from the external networks~~

We thank you again for the opportunity to provide comments to the Drafts and would appreciate the opportunity for further discussion and collaboration with NIST regarding the above.

Respectfully submitted,

Cricket Liu, EVP and Senior Fellow

Ed Hunter, CISO

Infoblox Inc.

November 6, 2023

[redacted]