GEORGETOWN UNIVERSITY

S^{2}ERC SECURITY AND SOFTWARE ENGINEERING RESEARCH CENTER

NIST Framework Response

The Cyber Threat Intelligence Information Sharing Exchange Ecosystem (CyberISE) program in the Security and Software Engineering Research Center (S^{2}ERC) at Georgetown[1] submits the following comments in response to the Request for Comments by the National Institute for Standards and Technology (NIST) on the *Preliminary Cybersecurity Framework*,[2] a NIST work product as set forth in Executive Order 13636, *Improving Critical Infrastructure Cybersecurity*.

The CyberISE program at Georgetown focuses on various research projects in conjunction with industry, standards organizations, and governments to enable automated cyber threat intelligence information sharing. The target for this work includes critical infrastructure sectors, as well as non-critical private enterprises, governments, and international organizations. Our research covers the technology, laws, regulations, and policies needed to make automated cyber threat intelligence information sharing a reality.

General Comments

When the CyberISE program started, we quickly learned there is no agreed ontology for cyber security. Even NIST-published glossaries, such as the NISTIR 7298,[3] have multiple, conflicting definitions for the same terms and leaves many important, yet

[1] See http://s2erc.georgetown.edu/projects/cyberISE/

[2] We will refer to the NIST document as the *Preliminary Framework* in this response.

[3] Kissel, R. (ed), *Glossary of Key Information Security Terms*, NISTIR 7298 Revision 2 http://nvlpubs.nist.gov/nistpubs/ir/2013/NIST.IR.7298r2.pdf, May 2013.

contentious terms undefined. Likewise, there is neither taxonomy of threat intelligence nor taxonomy of information sharing technology extant.

In order to foster meaningful cyber security standards, we need to know what we mean when we refer to a term. More importantly, since the attack indicators, that is the vectors, threats, motivations, and so on, are always changing, we need a taxonomy to be able to classify attacks we have not seen before so we can act on the attack based on its class (taxa).

Likewise, as the attacks and indicators are dynamic, it is important that we understand the classes of information that needs to be exchanged. Focusing on specific data items may not be productive, as items of interest may be different depending on the class and target of an attack. In fact, as an attack unfolds, the data that needs to be shared may evolve. Useful data elements to share also depend on the tools being used. Finally, the ability to share or the ability to use a particular data item often depends on how that data was acquired. That is, it is not necessarily dependent on what the data class is. These and other factors mean that on the one hand, there will be perpetual updates to tools and data to share, but on the other hand it means that to be effective, a lasting framework needs to be at a higher level than just prescribing elements to share or steps to take.

Something not mentioned at all in the *Preliminary Framework* is data obfuscation, data redaction, or data available through third-party security information providers. Although not stated directly by the *Preliminary Framework*, the information sharing mechanism reads as through it is a peer-to-peer, dumb receiver model, with no real capabilities at the receiver end to ingest and provide added value. In addition it does not leverage multiple security information providers, such that the receiver may build upon information by querying multiple providers and aggregating results. Some mention of these techniques, or at least acknowledging there are more options than just peer-to-peer, would be helpful.

At this point in time, the *Preliminary Framework* is not over prescriptive in general, but we do have some concerns that we will discuss in the narrative section of this response.

We are heartened to see the *Preliminary Framework* mention automated information exchange in Section C.2 and the need for a taxonomy in Section C.5. However, we believe there is a need to have more emphasis on the importance of having a standard taxonomy in the report.

Automated information sharing will serve to reduce cyber security risk. As such, we were pleased to see mention of Cybersecurity information sharing in the description of the different implementation tiers in the *Preliminary Framework*. We would not advocate for a one-size-fits-all prescription for participating in a particular exchange but it would be valuable to mention other levels of engagement, such as:

- Direct engagement (sharing and participating) with appropriate sector-specific Information Sharing and Analysis Center (ISAC)
- Engaging with US-CERT

- Participating in closed exchanges
- Subscribing to third-party feeds, such as from Symantec, IID, Verisign, etc.

Such engagement is hinted at around lines 386 – 389 and 419 – 421, but this could be more clear and expanded upon.

Let us return to the point that how an organization comes to hold a piece of Cybersecurity intelligence impacts to whom and under what circumstances the organization can share such information. Whether an organization shares information using automated tools or manually, the organization needs to be aware of these concerns. The concerns may be contractual, legal, or have liability issues. Nowhere does the framework core describe, for example, suitably identifying the classes of data and their respective controls or marking data elements with their provenance and permissions. We will identify a few candidate opportunities in our comments on the framework core for where the framework should point this out.

Where the *Preliminary Framework* does mention information sharing, it would also be helpful to layout some areas that need further study. For example, large dumps of information that do not pertain to a specific environment cause more work for the receiver to filter out or worse requires manual analysis to determine if a sub set of information pertains to their environment. During an attack, too much information can be damaging, as resources get expended evaluating unrelated information or completely ignoring important information because it's hidden within a large set of security information.

Finally, since we seem to enjoy maturity levels how about this as a proposal for information sharing maturity:

Level	Capability
0	No process for receiving, evaluating, or acting on third-party detection information. No feedback to sources.
1	Manual process for receiving, evaluating, and acting on third-party detection information. Minimal feedback to sources, if only a polite acknowledgement.
2	Single automated process for receiving, evaluating, and acting on third-party detection information. Feedback is primarily manual, but following a process.
3	Multiple means of automatically processing third-party detection information. Feedback is primarily manual, but following a formal process.
4	Continuous effort to seek out and integrate third-party sources of information. Some automated feedback to sources.

Narrative Comments

Supply Chain

Just as the framework applies to sectors and not just organizations, it will be important for the *Preliminary Framework* to have more emphasis on the supply chain. Line 244 implies the focus of the framework is within a given organization. ID.BE-1, 2, and 4 imply an organization may feed into other organizations and a failure in the organization in question could impact a critical infrastructure sector. However, what is not clear is that the organization in question can depend on other organizations, many of which may not be in an identified critical infrastructure / key resource (CIKR) sector. An organization may outsource some of its computing to the cloud. The fact the organization outsources computing may have little, no, or serious impact on the organization's ability to supply its part of the critical infrastructure. As an example, many hosted domain name service providers are not considered CIKR. However, if Web access to an organization's service is critical to the delivery of that service, and the organization has outsourced its DNS services, an attacker can attack the external DNS provider with the result being the same as if the organization itself was directly under attack.

We would not want the framework to ban outsourcing. In fact, using external service providers can significantly reduce an organization's risk exposure. For many sectors, third parties are much better able to protect these arcane enabling technologies. Thus, the framework should point out that when an organization evaluates its capabilities against the framework, the organization needs to include its entire supply chain, not just the critical infrastructure delivery supply chain, and consider options beyond the walls of the organization.

One possible reference for this could come from Section 2.5 of NIST SP 800-54 Rev. 4.

Information Sharing

The Respond function (lines 265ff) needs to mention information sharing. Our understanding is the whole point of the *Proposed Framework* is to protect critical infrastructure and key resources. That is a sector-wide aspiration, not an aspiration limited to any single organization. One known method of raising the posture of a sector is cyber threat intelligence sharing. The framework should make it a point of a responsive organization to include information sharing as part of its response.

The framework should mention cyber threat intelligence sharing as a source of detecting an existing or immanent breach in the Detect function (lines 259ff). In fact, it is hard to imagine how critical infrastructure organizations would learn about emergent risks (lines 419 – 421) without a robust cyber intelligence sharing program in place.

Core Comments

Identify

As discussed above, external service providers can have a critical impact upon the operation of an organization, even if that servic provider is not in a critical infrastructure sector or is not providing an identified key resource. The framework needs to specify Identifying and cataloging key external service providers and their service level agreements (SLAs) and the organization's mitigation strategy in the Identify section.

As an example, an organization may depend on non-redundant or moderately available (e.g., 99.9% SLA) cloud storage. Architecting the organization's IT process such that it leverages the cloud storage provider's geographic redundancy or by using a second, fully independent (including communication links) cloud storage provider may bring the cloud storage to an acceptable level. To reiterate, we would not want the framework to ban cloud storage, mandate multiple providers or technologies, or have other technological strictures. We do believe that an accounting of external providers and their impact on critical infrastructure delivery is important. ID.RA-3 looks to be focused internally on the organization. Perhaps it could be expanded to include external service providers. Alternatively, NIST could add a new subcategory covering external service providers.

One not well versed in risk analysis may read the Risk Assessment category and come away thinking that risk assessment is something an organization does once and then does not do again for a long time. ID.RA-2 hints that this is a perpetual process, as one expects to receive a constant stream of threat and vulnerability information. Conversely, ID.RA-3 appears to be a one-time audit. For the uninitiated, the framework should mention, in the framework introduction, that the framework represents a process, not a one-time audit.

Protective Technology

PR.IP-8 is an opportunity to highlight that information sharing is not an all-or-nothing exercise. Perhaps wording this subcategory as "Information sharing occurs with appropriate parties using appropriate data provenance and protection controls."

Note that the appropriate markings for automated information sharing are an area of active research underway in the CyberISE program at Georgetown. As such, it is much too soon to mandate a particular marking system, as the extant methods, such as Traffic Light Protocol, are being shown to not be sufficient to meet the legal and contractual obligations of critical infrastructure organizations. However, irrespective of the state of the technology, organizations need to be aware of and honor data protection obligations, which go beyond simply protecting PII.

Respond

As mentioned above, any sort of information sharing (e.g., the RS.CO subcategories) needs to honor limitations on sharing of particular data elements.

RS.IM should mention sharing lessons learned with others in the sector.

Privacy Issues

Identify:Governance

Data shared by partners may have stricter rules than the organization has on its own data. Thus, the organization needs to track, monitor, and honor such data sharing restrictions, more especially as the organization seeks to share data with external partners.

Likewise, the organization needs to ensure that its suppliers and service providers follow the appropriate data protection and disclosure rules that apply to the organization itself.

This governance section is entirely U.S.-focused. We appreciate the NIST effort is a U.S.-driven activity. However, many organizations in critical infrastructure sectors are multinational corporations. This makes it imperative for the organization to understand the various laws and rules governing personally identifiable information (PII) in the various jurisdictions the organization operates in. In addition, different states have different breach notification laws. Some jurisdictions have different rules for protecting suppliers' PII versus customers' PII. Therefore, this issue is not limited to multinational corporations.

Respond

In the Analysis category, the wording is not clear. It looks like the stricture is to make the policies accurate and complete, not the PII.

The Improvement category does not appear to say anything. Is there an action or consideration here? What is it?

Summary

We appreciate the open and consensus-driven process NIST established and is following for the creation of the Cybersecurity Framework. We would be happy to discuss our comments and answer any other questions you might have. Please contact the S^{2}ERC at Georgetown Center Director at eburger@cs.georgetown.edu or 202-687-4107.