Toward a Capability-Based Architecture for Cyberspace Defense

B. K. Done[a], K. D. Willett[b], D. W. Viel[c], G.W. Tally[c], D. F. Sterne[c], and B. Benjamin[c]

[a] US Department of Homeland Security

[b] US National Security Agency Information Assurance Directorate

[c] The Johns Hopkins University Applied Physics Laboratory

Abstract

Enterprises face huge challenges in protecting their networks from cyberspace attacks where the attacker has an asymmetric time advantage versus the defender's time to detect and respond. Attackers use tools to automate the attack process. Most security operations centers rely heavily on manual cyber defense operations to detect and respond to those attacks. Additionally, the lack of effective information sharing among enterprises enables attackers to reuse tools and techniques across multiple targets. The Department of Homeland Security and the National Security Agency Information Assurance Directorate, with The Johns Hopkins University Applied Physics Laboratory, are collaborating with industry and government to define a reference architecture that aims to reduce the cybersecurity event response time and limit the ability of attackers to effectively reuse tools and techniques. This joint effort is developing a concept for an active cyberspace defense ecosystem enabling near-real-time network defense at the enterprise level. The concept, called *Integrated Adaptive Cyber Defense* (IACD), supports cybersecurity operation automation, information sharing, and interoperability in the civilian Federal departments and agencies; state, local, tribal, and territorial governments; and critical infrastructure and key resources owners. The objectives of the IACD reference architecture are to (1) engage the cybersecurity industry in the refinement of the IACD vision and capabilities, (2) encourage and provide guidelines for implementing security automation and information sharing in enterprises with diverse architectures, and (3) promote commercial adoption of standardized machine-to-machine interfaces by communicating IACD needs and requirements to vendors.

This paper presents a high-level overview of the IACD reference architecture to inform, guide, and facilitate feedback from cyber service providers, network owners and operators, and product vendors on the capabilities and interfaces that can enable an agile, dynamically responsive, and resilient cyber infrastructure.

1 Introduction

Integrated Adaptive Cyber Defense (IACD) was initiated by the Department of Homeland Security (DHS) and the National Security Agency Information Assurance Directorate (NSA/IAD) in 2014 to effectively respond to two key issues that hamper the effectiveness of cyber defense for the enterprise. First is the adversaries' ability to reuse successful cyber-attack tools and techniques against multiple enterprises because of a lack of information sharing by those enterprises. Second, cyberspace-attack response times are currently too slow to address all cyber alerts, largely because existing response solutions require humans in the loop for multiple aspects of the four major elements of cyber event processing: sensing, sense-making, decision-making, and acting [1]. As a result, numerous cyber alerts are never analyzed due to lack of available analyst time, which enables attackers to penetrate more deeply before detection and remediation is achieved. The situation has become more severe because attackers have improved their efficient use of automation much more rapidly than the defenders [2], resulting in significant growth in the number and variety of attacks enterprises are experiencing, thus further taxing already overwhelmed analysts. Today's prevailing cyber defense approaches, which do not use efficient automation with continuous improvement, cannot be scaled to handle this load because it is labor intensive. Furthermore, most existing cyber defense technologies are not interoperable and are independently managed, which creates inefficiencies in the processing of cyber alerts. To reverse these trends, IACD seeks to improve cybersecurity automation and information sharing and encourage interoperability

between commercial tools for all Federal Departments and Agencies; state, local, tribal, and territorial governments; and critical infrastructure and key resources owners (hereafter referred to as the *IACD Community*). The concept relies on three foundational capabilities[1]: (1) *automation* enables automated sensing, sense-making, decision-making, and courses of action (COAs) (responses) [1] to provide network defense within cyber-relevant time; (2) *information sharing* enables rapid sharing of indicators, analytics, and effective COAs among enterprises and a coordinated response across portions of the IACD Community; and (3) *interoperability* enables a variety of commercial vendors' tools to function with each other without the need for pairwise, custom interfaces.

IACD integrates and extends results from several related DHS initiatives. The *Trusted Cyber Services* (TCS) [3] is composed of trust services, information and data management services, and analytics. TCS also provides reputation and enrichment services, shared cyber situational awareness, and integrated operational action services at a national level. *Trusted Information Sharing* (TIS) [4] enables controlled sharing of indicators of compromise, threat intelligence, analytics, and COAs across portions of the IACD Community. The *Enterprise Automated Security Environment* (EASE)[2] concept "envisions an environment in which automated and dynamic enterprise-level cyberspace defense capabilities—such as adaptive sensing, sense-making, decision-making, and acting—provide shared situational awareness and support response in cyber-relevant time." [5]

The IACD team is developing a capability-based reference architecture (RA) based on the Department of Defense (DoD) RA Description [6], which intends to be flexible, adaptable, and tailorable to the IACD Community. The IACD RA aligns with relevant Federal RAs, especially those focused on cybersecurity. IACD will support a *bring-your-own-enterprise* model and a vendor-agnostic plug-and-play operating environment to enable enterprises to select the commercial components and products that best suit their needs. The process of developing the RA will engage cyber service providers, network owners and operators, and product vendors to define the capabilities and interfaces that will enable an agile, dynamically responsive, and resilient cyber infrastructure.

2 Strategic Purpose

The strategic purpose of IACD is to improve the overall cyberspace defense of the IACD Community by enabling rapid information sharing among community members and automated response to cyber events using commercially available tools. Use of IACD will help remediate problems relating to (1) adversaries reusing similar tools and techniques against multiple targets and (2) an ongoing increase in the rate of cyber alerts and resulting incomplete processing of these alerts. The fundamental objective of IACD is "reducing response time from months to milliseconds" [1] so as to improve the response of enterprises to the cyber threat, and to be flexible and adaptable so that it can be tailored to the IACD Community.

The high-level goals for IACD to achieve this objective are to (1) improve the effectiveness of human capital for cyber operations within the IACD Community through *automation*; (2) degrade the attacker's ability to reuse their wares across the IACD Community through *enhanced information sharing*; and (3) remove barriers to adoption for the IACD Community through *interoperability*.

The IACD RA focuses on the capabilities required to achieve IACD goals and the integration of solutions that provide the capabilities. The IACD RA will (1) provide guidance for planning and implementing security automation and information sharing in enterprises with diverse information security infra-structures and (2) promote commercial adoption of standardized machine-to-machine interfaces by specifying IACD needs and requirements to vendors. IACD thus enables the establishment of enterprise communities able to share information and coordinate responses in near real time. Ultimately, enterprise owners and operators should be able to buy interoperable commercial products, thus allowing the use of multi-vendor solutions tailored to each enterprise. Enterprises will be able to supplement or replace

[1] A *capability* is an expression of a desired result agnostic of the solution that produces that result.
[2] EASE is in alignment with the NSA/IAD Adaptive Cyber Defense program and has similar goals.

vendor solutions over time while maintaining consistent operations. The IACD RA is not intended as a rigid specification, but supports a variety of solution implementations to suit diverse needs of a variety of enterprises.

IACD encompasses the active cyber defense capabilities that correspond to sensing and responding to cyberspace events for enterprises inside and outside of government. IACD seeks to be an organizing framework for cyberspace defenses that improves the current cyber system's organization and operations ability to respond to cyber events. As such, IACD may touch all aspects of the cyberspace and cybersecurity system, their component systems, and their operations. However, IACD does not specify particular implementations of solutions; rather, it describes how capabilities may interoperate together to create a more effective and scalable cyberspace defense.

The proposed IACD RA is influenced by a capability-based engineering approach [7], [8] driven by several motivating considerations: (1) avoid specifying the design and implementation of solutions to allow vendors to innovate; (2) encourage vendors to participate in a forum for the creation and sharing of open standards for interoperability; and (3) enable adoption by the IACD Community and managed security service providers by providing guidance that identifies required functionality and interface standards.

3 Principles

IACD follows a number of principles to further the goals enumerated in Section 2. Adoption of general, large-scale distributed systems engineering principles promotes robustness of the overall system and interoperability of components. Associated principles include the following:

a. *Essential Functionality*: For each capability, specify only the functionality necessary to ensure the capability meets the functional aims, including interoperability. This allows vendors to develop or modify their products in a way that requires minimal changes to any existing or planned products, thus raising the incentive for a vendor to participate in IACD.
b. *Loose Coupling*: Use asynchronous messaging instead of close-blocking remote service calls to reduce component interdependencies so they may operate independently, thus reducing the likelihood of multiple components becoming inoperable if a single component has problems.
c. *Statelessness*: Each component should execute its functionality without saving state between invocations to the extent possible. Being stateless increases the robustness of the solution and simplifies the protocol interactions among components.

To promote automation, IACD uses a variant of the Observe, Orient, Decide, and Act (OODA) control loop tailored to cybersecurity operations. These are designated as the sequence of Sensing, Sense-Making, Decision-Making, and Acting. IACD provides for machine implementations of these capabilities to facilitate the migration of people from in-the-loop to on-the-loop in cybersecurity operations. IACD facilitates adoption of technologies that scale with the size of the enterprise from a small group of a few hundred people to a combination of large departments with hundreds of thousands of people.

To promote sharing, IACD provides an efficient, robust, standards-based sharing capability to facilitate rapid sharing of threat information, analytics, and cyber event responses among the IACD Community.

To promote interoperability, IACD enables open-standards–based capability interfaces. Communications among the capabilities that comprise IACD will require solutions that implement messages compatible with each vendor's product. These open standards will support a market for security tools that emphasizes machine-to-machine information exchange and interoperability. The standards will also enable a *plug-and-play* approach to product integration. This allows components to be swapped in and out of an IACD system without reconfiguration of the whole. This reduces implementation timelines and cost and enables *dynamic composability* that invokes the most appropriate tool for the circumstances at hand.

IACD seeks to allow enterprises the opportunity to reuse as much of their existing cybersecurity infrastructure as possible and the freedom to specify and choose the implementations that comprise this infrastructure.

4 Technical Positions

The current IACD technical positions address standards and alignment to other key activities that provide or take influence from IACD.

4.1 Standards

IACD is a distributed enterprise system of systems. As such, it should be informed by and comply with industry standards to promote acceptance in the IACD Community and facilitate adoption and implementation by vendors. These standards include those for networking and messaging. In addition, vendors are encouraged to participate in developing the standards for the IACD messaging infrastructure and message set so that they are readily adopted by the IACD Community and industry.

In addition to general industry communications standards, several cybersecurity specifications are considered. These include the Structured Threat Information eXpression (STIX™) and Trusted Automated Exchange of Indicator Information (TAXII™) [9] specifications for sharing of cybersecurity indicators. The IACD Information Sharing capability provides a pathway for enterprises to leverage the specifications and capabilities being developed under the Federal Government's Enhanced Shared Situational Awareness (ESSA) [10] TIS and TCS initiatives. The ESSA TIS initiative is working to standardize the sharing of COAs, analytics, and indicators with other ESSA member enterprises, while still supporting proprietary interfaces to leverage critical external information sources. ESSA uses the STIX and TAXII protocols for information sharing, and efforts are underway to define data exchange protocols for COAs. Additionally, the TCS initiatives, beyond ESSA, provide new capabilities such as the Cyber Weather Map, reputation services, and enrichment services. The IACD team is working with stakeholders to define the data exchange standards and protocols required for these services.

4.2 Relationship to Other Frameworks and Initiatives

IACD does not live in isolation. Many other efforts, systems, and environments exist with which IACD should conform; four in particular are described next.

The National Institute of Standards and Technology (NIST) Cybersecurity Framework [11] provides guidance for implementing cybersecurity functions spanning Identify, Protect, Detect, Respond, and Recover and does so in a way that does not prescribe a specific solution for all organizations. IACD is in alignment with these interests, including being relevant to all framework functions. Both the NIST Framework and IACD provide implementation flexibility that enables vendors to continue to develop more effective cyber capabilities while achieving the required functionality and interoperability. IACD maps well to automation of the NIST Cybersecurity Framework's Detect function and is addressed by several capabilities including Sensors and Actuators, Sensor/Actuator Control and Data Normalization, and Sense-Making. Decision-Making, Response Controlling and Actuators support the Respond and Recover functions. Additionally, IACD supports information sharing and collaboration not specifically found in the NIST Cybersecurity Framework, but critical to more effective community-wide cybersecurity capabilities.

The Cyber Kill Chain defines a series of steps an adversary takes to penetrate and establish a persistent presence in the target system. This pattern of action is similar to that which the military uses for engaging an enemy. The steps of the Cyber Kill Chain include reconnaissance, weaponization, delivery, exploitation, installation, command and control, and actions on targets. The kill chain analogy shows that the attacker must complete each step in turn for the attack to succeed. Preventing any step disables the entire attack. Currently, defenders too often detect an attack during the final stage of the kill chain, after the damage is done [12]. Mapping adversary kill chain indicators to defender COAs, identifying patterns

that link individual intrusions into broader campaigns, and understanding the iterative nature of intelligence gathering allows construction of an effective defense. [13] IACD will enable identification of attacks before the "action" stage of the kill chain, thereby enabling early response to the attack to minimize or prevent damage.

The Cyber Security Information Act promotes the sharing of cybersecurity information among enterprises, with which the goals of IACD are in close alignment. IACD enables the kind of sharing of cyber information prescribed by the Act, and promotes the use of sharing mechanisms in the government and private sectors. Furthermore, the automation goals of IACD enable a greater level and more efficient sharing than is possible with current cybersecurity system deployments.

The Cybersecurity Strategy and Implementation Plan (CSIP) for the Federal Civilian Government [14] identifies goals for cybersecurity that overlap those of IACD in many areas. In particular, CSIP calls for (1) timely detection and rapid response to cyber incidents, (2) rapid recovery from incidents and accelerated adoption of lessons learned, and (3) efficient and effective acquisition and deployment of existing and emerging technology.

5 Patterns

Figure 1 describes the top-level IACD capabilities and functions, integrated through a local Control Message Infrastructure that provides timely delivery of capability outputs to one or more recipients. The sensors, actuators, and external data feeds shown in the figure are external to the IACD RA. The sensors sense cyber events within an enterprise, and the actuators implement response actions within an enterprise. The grey arc in Figure 1 shows the flow of the OODA loop, which directs the sequence of Sensing, Sense Making, Decision Making, and Acting in an IACD system.

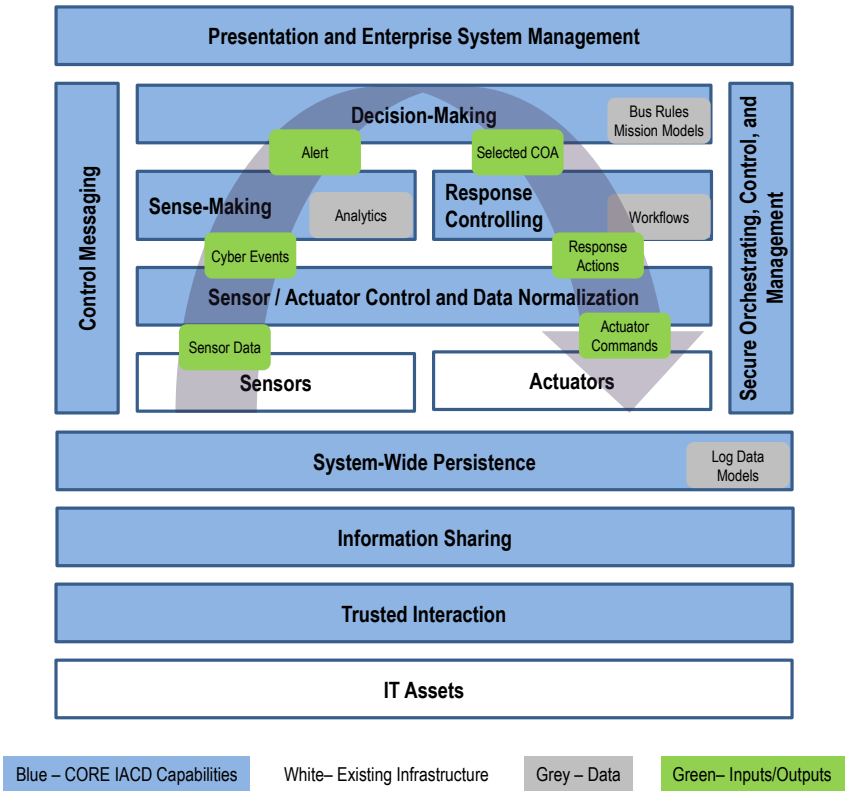

Figure 1 Integration of IACD Capabilities through the OODA Loop

Notional top-level IACD capabilities are as follows:

a. *Secure Orchestration, Control, and Management*: Encompasses coordinating, controlling, and managing interactions among the IACD capabilities previously described. Provides control of the OODA loop execution. Includes coordinating IACD service automation and information sharing; collecting associated operational data; starting, stopping, directing, and reconfiguring IACD components; monitoring IACD components for operational activity anomalies; and maintaining IACD configuration information. These activities may be thought of as a distinct and potentially centralized capability, as suggested by Figure 1, or as aspects of (that belong to) other capabilities described next. The latter view is perhaps more naturally aligned with decentralized solution architectures, in which these activities are intrinsically coordinated behaviors that arise from cooperative interactions among other capabilities.
b. *Control Messaging*: Provides secure, timely, and loosely coupled message-based communication functionality among the IACD capabilities within an enterprise. Provides a standard set of messages for compliant components to implement. This capability facilitates passing commands and data between components produced by a variety of commercial vendors.

    1. Supports the publish/subscribe paradigm, which decouples data producers and consumers.
    2. Supports loose coupling of products from diverse vendors that may not have previously integrated their products. Iterations of the IACD RA will continue to further specify the content and format of the data in compliance with standards being developed by industry, but may support multiple transport protocols for Control Messaging.
    3. Some aspects may be implemented by a message bus, with different buses being appropriate to address the security, scalability, resiliency, and topology requirements of different enterprises.

c. *Sensor/Actuator Control and Data Normalization*: Enables secure communication of data, commands, and status with heterogeneous collections of sensors and actuators, many of which may use proprietary protocols; and translates, normalizes, and characterizes sensor and actuator messages to or from other IACD capabilities. Abstracts the specific sensors and actuators interfaces for other IACD capabilities to facilitate generalized interfaces. Inputs: Sensor Data and Responses. Outputs: Cyber Events and Actuator Commands.
d. *Sense-Making*: Evaluates cyber events and intelligence data in the context of the enterprise to determine whether a cyber alert should be raised, and characterizes the nature of the alert; and provides human operators with a visualization of the current security state. This capability includes analytics to identify potentially malicious behavior based on all forms of ingested sensor inputs. This capability identifies the target(s) of the event, the type of event, and typically includes a certainty assertion. Some analytics may involve human analysis of visual output. This capability also provides mechanisms to allow analytics to be managed and tested in a safe environment. Inputs: Cyber Events. Outputs: Cyber Alerts.
e. *Decision-Making*: Given a cyber alert generated by Sense-Making, recommends an appropriate response based on enterprise policies and risk and impact to the enterprise; includes human-in-the-loop review, as necessary. This requires identifying and evaluating alternative COAs for cyber alerts, which in turn requires understanding the mission impact of the events and responses. Inputs: Cyber Alerts. Outputs: Selected COAs.

    1. Enables the specification of mission priorities and mission dependencies on cyber capabilities.
    2. Enables the creation, editing, and management of preplanned COAs, which may include requesting more data (enrichment), directing countermeasures, sharing information with other community members, notifying users, etc.

3. When selecting a COA, controls the level of human involvement within the COA to approve response actions. Provides an interface for privileged human users to select or modify automatically recommended COAs.
4. Enables a privileged user to analyze shared COAs received from a community member and selectively adopt, delete, replace, or augment their individual steps to make the resulting COA consistent with the recipient enterprise's COA policies and practices.
5. Supports evaluation of candidate COAs in a safe environment to determine effectiveness and impact.
6. Enables the specification and enforcement of enterprise-specific COA policies ("sanity checking") that constrain the implementation of COAs to reduce the likelihood of COA errors and collateral damage.
7. Manages and provides persistent storage for mission models, COAs, and COA policy used during the Decision-Making process.

f. *Response Controlling*: Supervises the execution of response actions prescribed by Decision-Making; sequences workflows; coordinates responses among multiple organizations; and provides status to other capabilities of pending, ongoing, and completed COAs. Inputs: Selected COAs. Outputs: Response Commands.
g. *System-wide Persistence*: Provides the ability to store, organize, and make accessible a variety of sensor data, intelligence, and enterprise models including the following:

1. Log Data include all forms of sensor data with associated metadata.
2. Intelligence Repository includes archived intelligence data (e.g., threats, vulnerability, and indicators) from multiple sources with the associated metadata.
3. IACD Configuration Information describes the IACD constituent components and their configurations. Describes (abstractly) the set of messages each sensor or actuator produces and consumes, e.g., "Block IP Address X."
4. Network, Host, Software Configuration Information describes network device, host, and software configurations on defended networks. Identifies service dependencies on other cyber assets; *these enable mission dependency modeling and mission impact assessments.*
5. Network Models describe network topology and all network elements (hardware and software, blue and grey).
6. Performance Measurements record metric data for performance measurement and analysis.

h. *Information Sharing*: Enables secure communications with (1) other IACD Community members for IACD-standardized exchange of indicators of compromise, analytics, recommended COAs, cyber event data, and requests for coordinated action; and (2) non-IACD sources to obtain reputation services, intelligence feeds, and signature updates using protocols specific to each source or recipient.
i. *Trusted Interaction*: Preserves the confidentiality, integrity, and availability of IACD and IACD-managed systems. This capability provides authentication and access controls; manages access and permissions; protects passwords; and protects classification levels, compartments, applications, and services. It also ensures network monitoring and data collection and storage are in accordance with privacy laws and requirements.
j. *Presentation and Enterprise System Management*: Manages the IACD configuration and monitors system health.

6 Conclusion

IACD is an ongoing initiative to improve cyber defense capabilities for government and commercial organizations. The near-term steps for the effort include working with prospective enterprise adopters and the vendors to solicit feedback on the IACD RA and to facilitate vendor discussion and agreement on

interoperability and interface standards. Tasks for these steps may include supporting working groups for the vendors to adopt standard data formats, protocols, and interoperability specifications; conducting proof-of-concept tests with multiple organizations to evaluate the architecture; and evaluating the results and adjusting the architecture as required to meet the IACD objective.

The IACD activity intends to reduce cyber event response time from months to milliseconds and limit the ability of attackers to successfully reuse tools and techniques, thus increasing the cost of attack. This paper provides an overview of the IACD capability-based RA. The architecture supports IACD goals by (1) encouraging and providing guidelines for implementing security automation and information sharing in enterprises with diverse legacy architectures and (2) promoting commercial adoption of standardized machine-to-machine interfaces by communicating IACD needs and requirements to vendors.

The IACD effort is moving forward on engaging users and vendors. IACD has laid out the vision. It is now up to these groups to advance this vision of a more secure world by participating in the crafting of standards and specification and in pilot projects to test and field the resulting technology. Interested parties should contact one of the authors for more information.

7 References

[1] P. M. Fonash and P. Schneck, "From Months to Milliseconds," *Computer Magazine*, IEEE Computer Society, January 2015.

[2] Verizon, "Data Breach Investigation Report," 2015. [Online]. Available: http://www.verizonenterprise.com/DBIR/2015/.

[3] P. Schneck, "Modern Department of Homeland Security Cyber: Our Vision Forward," *RSA Conference*, San Francisco, 2015.

[4] US Department of Homeland Security, "Information Sharing Architecture (ISA)," 19 December 2013. [Online]. Available: http://csrc.nist.gov/groups/SMA/ispab/documents/minutes/2013-12/essa_isa_intro_requirements_overview.pdf.

[5] US Department of Homeland Security, "Enterprise Automated Security Environment (EASE) Request for Information (RFI)," Solicitation Number," 2014.

[6] "DoD Reference Architecture Description," [Online]. Available: http://dodcio.defense.gov/Portals/0/Documents/DIEA/Ref_Archi_Description_Final_v1_18Jun10.pdf.

[7] R. S. Swarz and J. K. DeRosa, "A framework for enterprise systems engineering processes," *Proceedings of International Conference on Software and Systems Engineering*, 2006.

[8] K. D. Willett, "Capability-Based Engineering Approach to Integrated Adaptive Cyberspace Defense (IACD)," *Information Assurance Directorate Symposium*, 2015.

[9] MITRE, "Support for STIX/TAXII," 19 August 2015. [Online]. Available: http://stixproject.github.io/supporters/.

[10] NIST, "Information Sharing Architecture," 19 December 2013. [Online]. Available: http://csrc.nist.gov/groups/SMA/ispab/documents/minutes/2013-12/essa_isa_intro_requirements_overview.pdf.

[11] NIST, "Framework for Improving Critical Infrastructure Cybersecurity," Version 1.0, 2014.

[12] The Johns Hopkins Applied Physics Laboratory, "Integrated Adaptive Cyber Defense: Spiral 2," 2015.

[13] E. M. Hutchins, M. J. Cloppert and R. M. Amin, "Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains," *Leading Issues in Information Warfare & Security Research*, Vol. 1, p. 80, 2011.

[14] Office of Management and Budget, Executive Office of the President of the United States, "Cybersecurity Strategy and Implementation Plan (CSIP) for the Federal Civilian Government," 2015.

[15] US-CERT, "Cyber Glossary," US-CERT, [Online]. Available: https://niccs.us-cert.gov/glossary.