IDA | RESEARCH NOTES

CHALLENGES IN CYBERSPACE: STRATEGY AND OPERATIONAL CONCEPTS

November 2019

IDA is the Institute for Defense Analyses, a nonprofit corporation operating in the public interest.

IDA's three Federally Funded Research and Development Centers answer the most challenging U.S. security and science policy questions with objective analysis leveraging extraordinary scientific, technical, and analytic expertise.

The summaries in this edition of *IDA Research Notes* were written by researchers within the following three IDA research groups. The directors of those divisions would be glad to respond to questions about the specific research topics or related issues.

Information Technology and Systems Division (ITSD), Dr. Margaret E. Myers, Director (703.578.2782, mmyers@ida.org)

Operational Evaluation Division (OED), Mr. Robert R. Soule, Director (703.845.2482, rsoule@ida.org)

Strategy, Forces and Resources Division (SFRD), Mr. John C. Harvey, Jr., Director (703.575.4530, jharvey@ida.org)

IDA

Institute for Defense Analyses
4850 Mark Center Drive
Alexandria, Virginia 22311
ida.org

Challenges in Cyberspace: Strategy and Operational Concepts

The first *IDA Research Notes on Challenges in Cyberspace* was published in 2011, not long after the Department of Defense recognized cyberspace as the fifth operating domain.

In the keynote article of that publication, retired **General Larry Welch** described ways to mature our understanding of cyberspace operations. Cyberspace is a domain, he wrote—a place, not a mission. As in the other more commonly understood domains—land, sea, air, and space—military superiority is derived from our freedom of action in, through, and from cyberspace, and from our ability to deny adversaries freedom of action at the times and places of our choosing. General Welch further discussed the activities that produce the desired military effects across the spectrum of cyber operations: constructing cyberspace, passive defense, active defense, exploitation or operational preparation of the environment, and attack.

The most recent *IDA Research Notes on Challenges in Cyberspace*, published in 2018, presented multiple aspects of IDA research related to the human dimension of challenges in cyberspace. The articles in this issue of *IDA Research Notes* discuss our research related to strategy and operational concepts for addressing challenges in cyberspace.

In the opening article here, **Dr. Michael Fischerkeller** highlights the positive progress in U.S. cyber strategy maturation over the past year but cautions that, going forward, policymakers must be disciplined in their discussions and expectations of the new cyber strategic framework, comprising strategies of defense, deterrence, and persistent engagement. This caution is borne of numerous examples of the defense community repeatedly overlaying strategic concepts and ends developed for one environment atop another without a careful consideration of whether they are in strategic alignment. Dr. Fischerkeller specifically discusses deterrence by denial in cyberspace and equating the generation of a deterrence effect with security.

Moving from the strategic to the operational, **Erick McCroskey** and **Charles Mock** show how they have enhanced well-established joint military symbology to help cyber warriors easily express operational concepts in cyberspace. Their cyberspace operational graphics will allow cyber planners and operators to convey mission-relevant information to warfighters who are unfamiliar with the technical details of cyberspace. Leveraging warfighter familiarity with common symbology language will enhance rapid understanding and decision-making and provide the joint commander coherent operational doctrine and accompanying graphics that will enable him to understand, plan, and fight the cyber battle.

The next article, by **Dr. Allison Goodman**, describes how we support the congressionally mandated Cybersecurity Assessment Program on behalf of the Director of Operational Test and Evaluation (DOT&E) and the

analytical methods used to conduct these data-based evaluations. The IDA research in support of DOT&E produces recommendations for both local and department-wide defensive approaches and vulnerability mitigation strategies.

The fourth article, by **Dr. Vik Kulkarni** and **Dr. Shawn Whetstone**, continues the discussion of the Cybersecurity Assessment Program with a summary of their analysis of Combatant Command training exercises for fiscal years 2014 through 2016. Using data collected on attack detection, defensive responses, and operational effects, they developed an analytical framework for operational cybersecurity assessments and used that framework to inform defensive strategy principles and provide recommendations for improving detection of attacks.

In the concluding article, **Dr. David Alberts** discusses the Office of Personnel Management data breaches and the DoD and IDA roles and activities in responding to the breaches. He explains why an agility-driven transformation of the security clearance vetting ecosystem is the only logical option for the future.

The Structural and Strategic Imperative: The Need for Persistent Engagement

Michael P. Fischerkeller

The Challenge: The emergent strategic framework for cyberspace, comprising strategies of defense, deterrence, and persistent engagement, is well-suited for the cyber strategic environment but requires vigilance by policymakers not to succumb to strategic conflation and confusion.

It is critically important to avoid adopting strategic concepts developed for one environment and indiscriminately overlaying them in an attempt to secure another.

Background

In May 2017, Richard Harknett and I argued that a strategy of deterrence should not be the central strategy for securing and advancing U.S. national interests in, through, and from cyberspace due to the unique structural and operational characteristics of cyberspace (Fischerkeller and Harknett 2017, 381–93).[1] In the 18 months since, dramatic changes in U.S strategic guidance—the National Security Strategy and National Defense Strategy and, most recently, the Department of Defense Cyber Strategy—now reflect this perspective, arguing that a central challenge to the U.S. lies in strategic competition short of armed conflict in cyberspace (as well as other domains) and that a strategy of deterrence is not an effective anchoring approach for ensuring security in this strategic competitive space (Department of Defense 2018a, 2; Department of Defense 2018b, 2; White House 2017, 3, 31). We argued in 2017 that recognizing the emergence of this new strategic competitive space, one in which adversaries are realizing strategic gains through cumulative effects from cyber campaigns, was a necessary first hurdle to overcome in developing an effective strategic approach to the same.

In February 2018, U.S. Cyber Command published its *Command Vision*, a document describing the cyberspace strategic environment as a new strategic competitive space below armed conflict and prescribing a strategic approach for securing U.S. national interests in cyberspace consistent with what we argued in 2017: a strategy of persistent engagement (United States Cyber Command 2018). This was captured in guidance in the recently released Department of Defense Cyber Strategy, which argues the Department must preserve

[1] Dr. Harknett and I have written several other IDA monographs on cyber strategy since this publication (and have a few more in the pipeline). I'm grateful for his review of and thoughtful recommendations for this article.

U.S. military advantages and defend U.S. interests in, through, and from cyberspace by taking action to contest persistently adversaries' malicious cyber activity during day-to-day competition (Department of Defense 2018b, 4). The strategy rests on recognizing the distinctive structural features of cyberspace and how adversaries have been exploiting them; it is, thus, not a choice if the United States is seeking security in this space, but a structurally and strategically driven imperative to reorient our approach to security. We argued in 2017 that cyberspace's structural feature of interconnectedness and its core condition of constant contact create a strategic necessity to operate persistently in cyberspace. The states that master persistent engagement will not only be more secure in cyberspace, they will also position themselves to enhance their national power relative to others.

From our perspective, then, the United States has made significant progress in the past year toward arriving at a strategic approach that is aligned more effectively with the unique structural and operational characteristics of cyberspace. We are cautiously optimistic that strategic development will continue along a path more consistent with what we've argued is necessary to ensure U.S. security. We use the modifier "cautiously" because a strategy of deterrence and the strategic concepts associated therewith continue to hold a strong place in the minds of U.S. policymakers and other senior leaders. This is to be expected given that a strategy of deterrence has been the dominant U.S. security strategy for the last 70 years. And so the defense community must now be vigilant in ensuring that a strategy of deterrence and a strategy of persistent engagement are viewed as two distinct, yet complementary, strategic approaches grounded in and developed for uniquely different strategic environments. It is critically important to avoid adopting strategic concepts developed for one environment and indiscriminately overlaying them in an attempt to secure another. By understanding the differences between the two strategies, we, in the end, can develop an overarching framework that leverages both (properly applied) to advance U.S. interests in cyberspace. In support of that objective, this article highlights two often seen or heard examples of how the discourse of deterrence too easily, and potentially dangerously, continues to bleed into discussions of a strategy of persistent engagement: the strategic concept of deterrence by denial and equating the generation of a deterrence effect with security.

Deterrence by Denial

When discussed in U.S. strategic guidance, *deterrence is often split into two forms: deterrence by punishment and deterrence by denial.*[2] Given the infrequency of cyber operations targeting the United States having caused effects equivalent with

[2] This is often referred to as *deterrence by cost imposition*, a choice not made here because it will be argued that deterrence both by punishment and by denial impose costs.

armed attack, one could argue that a strategy of deterrence by punishment has been effective in the strategic space of armed conflict (i.e., the strategic space above the threshold of armed attack). It is likely states and other actors understand that such operations would justify under U.N. Article 51 a conventional force (or even nuclear) response in self-defense. The potential (or promise) to inflict punishment on adversaries considering cyber operations equivalent with armed attack, however, has done nothing to reduce the scale or scope of strategic cyber campaigns and operations generating effects short of that threshold.

In fact, relative U.S. restraint in cyberspace appears to have incentivized adversaries' pursuit of strategic objectives in the cyber strategic competitive space short of armed conflict precisely to avoid a conventional confrontation with the United States (Fischerkeller and Harknett 2018). Interestingly, war (or pushing the threshold of war), we've argued, would, in almost all cases, actually represent a failure of an adversary's cyber strategy. Ten years of cyber campaigns and operations by the same have demonstrated it is possible to generate strategic effects through cyber campaigns/operations short of armed attack.

In efforts to mitigate the consequences of adversary cyber operations short of armed attack equivalence, U.S. policymakers have doubled down on the strategic concept of deterrence by denial, routinely identifying hardening and resiliency as operational capabilities/activities necessary to support the same (Department of Defense 2017, 3, 7; United States Congress 2017, 696; United States Congress 2018, 491).[3] Given the unique characteristics of cyberspace that Harknett and I have discussed, we argue that the primary function of these capabilities/activities actually supports a strategy of defense in the cyberspace strategic environment, not a strategy of deterrence by denial.

Glenn Snyder was the first scholar to make a sharp distinction between the objectives of the strategic concepts of deterrence by punishment and denial and a strategy of defense and did so within the context of the nuclear strategic environment (Snyder 1958; Snyder 1961).[4] Snyder argued that an overall strategy of deterrence—"by denial" coupled with "by punishment"—targets an adversary's *intentions* to attack, whereas a strategy of defense reduces an adversary's *capability* to damage or deprive the defender if attacked. Thus, the "deterrent value" of military capabilities is their effect in reducing the likelihood of adversary military operations, and the "defense value" of military capabilities is their effect in mitigating the adverse consequences

[3] The Defense Science Board stated that "steps to promote deterrence by denial include improving cyber defenses and increasing resilience of key systems to attack" and "hardening and increasing the resilience of the most vital critical infrastructure systems – including electricity, water, and waste water – is urgently needed to bolster deterrence by denial."

[4] In 1959, Glenn Snyder published a foundational manuscript describing the difference between deterrence by punishment and deterrence by denial. Snyder incorporated much from that manuscript into a follow-on publication discussing a strategy of defense.

of adversary operations, including losses of territory or war damage.

John Mearsheimer, often cited as one of the principal architects of conventional deterrence theory, argues that conventional deterrence is primarily based on deterrence by denial, described as the ability to prevent an adversary from achieving its objectives through conflict (Mearsheimer 1983, 24, 30, 64, 206–08; see also Lebow 1981, 248; Rhodes 2000, 243–47; Van Evera 1999, 30–4). If states typically seek short and low-cost conflicts, he argued, then conventional deterrence largely depends on convincing an adversary that it cannot achieve its objectives rapidly or efficiently. In this context, the deterrent effect is achieved in large part by the possibility of getting bogged down in a long and costly war of attrition.

According to Mearsheimer, "… deterrence is best served when the attacker believes that his only alternative is a protracted war: The threat of a war of attrition is the bedrock of conventional deterrence" (Mearsheimer 1983, 206–7). That is, deterrence by denial seeks to make aggression unprofitable by rendering the target harder to take, harder to keep, or both (Mitchell 2015). A credible deterrence-by-denial posture, then, requires that a would-be attacker believe a defender's forces are inherently designed to fight and win a conflict in the event of a deterrence failure (Gerson 2009). The defender's capabilities have to be perceived as lethal and able to inflict substantial pain (Mitchell 2015). Like deterrence by punishment, deterrence by denial also threatens to impose costs but differs from punishment in that it threatens to impose costs during the act of aggression and in the place that it occurs, as opposed to at a time and place of the deterrer's choosing (Mitchell 2015).

To come full circle, policymakers continue to cite hardening and resilience as operational capabilities/activities that support a strategy of deterrence by denial in cyberspace. When considering foundational conventional deterrence scholarship, however, those operational capabilities/activities clearly do not match the requirements for an effective deterrence by denial strategy (i.e., hardening and resilience do not represent a lethal posture populated by capabilities that would support an attrition-based response to a cyberattack).

Moreover, the viability of any strategy relying on the concept of attrition warfare in cyberspace finds little support in the scholarly literature (see Brantly 2015).[5] Again, this is structurally determined – the features of cyberspace do not succumb to attrition. This is not to argue that investments in hardening and resilience are not worthwhile. They certainly are, because they support the critical objective of a strategy of defense as described by Snyder, which, to reiterate, is to reduce an adversary's capability to damage or deprive the defender.

[5] Brantly argues that the virtual nature of the cyberspace domain makes attrition warfare an inefficient allocation of resources and unlikely to achieve sustained effects on agile targets.

Why This Matters

If policymakers believe that investments in hardening and resilience are supporting a U.S. strategy of deterrence by denial in cyberspace, they will expect to see results toward that end. Continuing and intensifying adversary activity against U.S. interests accessible in, through, and from cyberspace strongly suggest that hardening and resilience are not affecting our adversaries' *intentions* to attack. There is little doubt that hardening and resilience capabilities/activities are reducing adversaries' *capability* to damage or deprive the United States (i.e., supporting a strategy of defense), but policymakers should not hold false expectations that continuing to invest in these capabilities/activities will lead to changes in our adversaries' intentions to continue to attack in, through, and from cyberspace.

Deterrence Does Not Necessarily Equal Security

In the nuclear and conventional strategic environments, the central strategy for the United States for the last 70 years has been a strategy of deterrence. In both environments, the generation of a deterrence effect has been equated with security. As was argued above, when applied to the cyberspace strategic environment over the past several years, the threat of punishment appears to have deterred adversary cyberattacks equivalent with armed attack, but the overall strategic end for the United States has, nonetheless, been a gradual erosion of its power relative to adversaries over the same period.

This erosion is in many ways directly attributable to strategic gains accumulated by adversary cyber campaigns/operations short of armed attack, gains resulting from cyber campaigns or operations targeting intellectual property, military overmatch, and social cohesion. It would be dubious, then, to associate the generation of a deterrence effect in cyberspace with the realization of security in the same. And yet, early commentary on the 2018 DoD Cyber Strategy's discussion of persistent contestation argues that a measure of effectiveness of a strategy of persistent engagement (or, at least, the "defend forward" aspect of it) should be its "deterrence value" (Pomerlou 2018).

This is another example of the indiscriminate overlaying of a strategic concept associated with a strategy of deterrence (i.e., a deterrence effect) onto the cyber strategic environment. The cyber strategic competitive space below the threshold of armed conflict is characterized by constant contact manifesting in continuous engagement, enduring campaigns, and competitive interactions between friends and foes alike. As Harknett and I have argued, operational persistence is a strategic imperative for states seeking security in this space (Fischerkeller and Harknett 2017; 2018). Any strategy for cyberspace that intends to coerce comprehensive operational restraint on the part of an adversary (i.e., generate a deterrence effect) will fail because such an outcome is not aligned with the strategic imperative presented by cyberspace.

Persistent engagement accepts this imperative and, consequently, does not claim deterrence as a strategic objective. Rather we have argued that persistent engagement has an operational objective of inhibiting, not eliminating, adversary cyber campaigns or operations that threaten U.S. national interests. Additionally, it has the strategic objectives of preventing extended or enduring imbalances (negative or positive) from emerging in the cyber strategic competition below the threshold of armed conflict and supporting the development of mutual understandings with adversaries of acceptable/unacceptable behavior in the same (Fischerkeller and Harknett 2018). It is through the combination of these effects that security can be realized in this strategic competitive space.

Conclusion

A significant reorientation has occurred in U.S. strategic guidance published over the past 12 months. Strategic competition below the threshold of armed conflict is now considered a central challenge to U.S. national security, and the cyberspace strategic environment has been accepted as a primary space in which this competition occurs. This reorientation was appropriately accompanied by an adjustment to the U.S. strategic framework, one that now emphasizes the equal importance of a strategy of deterrence and a strategy of persistent engagement in ensuring national security short of armed conflict.

This article encourages strategic discipline and vigilance on the part of the defense community to hold fast and true to understandings of the structural foundations and strategic features of each so that the strategic concepts and intended effects associated therewith are not conflated, confused, or otherwise used indiscriminately. Such muddling could result in a range of inefficient or even potentially dangerous behaviors, ranging from a misallocation of scarce resources to a misconception of intended strategic effects, none of which would well serve the pursuit of U.S. national security.

References

Brantly, Aaron F. 2015. "Strategic Cyber Maneuver," *Small Wars Journal*. http://smallwarsjournal.com/jrnl/art/strategic-cyber-maneuver.

Department of Defense. 2017. *Department of Defense – Defense Science Board Task Force on Cyber Deterrence*. Washington, DC: Department of Defense.

Department of Defense. 2018a. *Summary of the 2018 National Defense Strategy of the United States of America*. Washington, DC: Department of Defense.

Department of Defense. 2018b. *Summary: Department of Defense Cyber Strategy*. Washington, DC: Department of Defense.

Fischerkeller, Michael and Richard J. Harknett. 2017. "Deterrence is Not a Credible Strategy for Cyberspace," *Orbis* 61(Summer): 3.

Fischerkeller, Michael and Richard J. Harknett. 2018. *Persistent Engagement, Agreed Competition, Cyber Interactions Dynamics, and Escalation*. Alexandria, VA: Institute for Defense Analyses.

Gerson, Michael S. 2009. "Conventional Deterrence in the Second Nuclear Age," *Parameters: United States Army War College Quarterly*.

Lebow, Richard Ned. 1981. *Between Peace and War: The Nature of International Crisis*. Baltimore, MD: Johns Hopkins University.

Mearsheimer, John J. 1983. *Conventional Deterrence*. Ithaca, NY: Cornell University Press.

Mitchell, A. Weiss. 2015, August 12. "The Case for Deterrence by Denial." *The American Interest*. https://www.the-american-interest.com/2015/08/12/the-case-for-deterrence-by-denial/.

Pomerlou, Mark. 2018, September 28. "It's a New Era for Cyber Operations, but Questions Remain." *Fifth Domain*. https://www.fifthdomain.com/dod/2018/09/28/its-a-new-era-for-cyber-operations-but-questions-remain/.

Rhodes, Edward. 2000. "Conventional Deterrence," *Comparative Strategy* 19: 243–47.

Snyder, Glenn H. 1958. *Deterrence by Denial and Punishment: Research Monograph No. 1*, Princeton University.

Snyder, Glenn H. 1961. *Deterrence and Defense: Toward a Theory of National Security*, Princeton University Press: Princeton, NJ.

United States Congress. 2017. *National Defense Authorization Act 2018*. S. 1519, Report No. 115-125, Section 1621, (c) Denial Options.

United States Congress. 2018. *John S. McCain National Defense Authorization Act 2019*, H.R. 5515, Section 1636, (c) Denial Options.

United States Cyber Command. 2018. *Command Vision for U.S. Cyber Command: Achieve and Maintain Cyberspace Superiority*. Washington, DC: Department of Defense.

Van Evera, Stephen. 1999. *Causes of War: Power and the Roots of Conflict*. Ithaca, NY: Cornell University Press.

White House. 2017, December. *National Security Strategy of the United States of America*. Washington, DC: White House.

Dr. Fischerkeller is a Research Staff Member in IDA's Information Technology and Systems Division. He holds a Doctor of Philosophy in international security from the Ohio State University.

Operational Graphics for Cyberspace

Erick D. McCroskey and Charles A. Mock

The Challenge: Specialized sets of symbols that convey information and understanding faster than text alone have been part of military tactics, strategy, and the operational art since armies became too large for personal observation on the battlefield. The Department of Defense (DoD) established cyberspace as the newest warfighting domain via doctrinal guidance in 2011, yet cyber warriors still lack a coherent set of symbols that allow them to convey the intricacies of cyber warfare to the joint warfighting community. The inability of cyber warriors to easily express operational concepts inhibits identification of cyber key terrain, the development of tactics and strategies, and the execution of command and control.

Cyber organizations do not lack for symbols and graphics—network diagrams are ubiquitous—but these symbols do not conform to joint warfighting doctrine.

Introduction

A sergeant looks at an arrow marked in grease pencil on a laminated map and knows that a machinegun position lies ahead. The large projection screen showing a map with a blue rectangle encompassing an oval gives the Joint Task Force commander assurance that a tank battalion defends key terrain. A picture is worth a thousand words.

The primitive state of cyber operational graphics, and the resulting lack of effective communication between cyber and physical domain warriors, deemphasizes operational campaign design and the application of the principles of war in cyber operations. This increases the likelihood that physical domain warfighters will accept dangerous risks because they have little conception of what is really happening on their networks.

Cyber organizations do not lack for symbols and graphics—network diagrams are ubiquitous—but these symbols do not conform to joint warfighting doctrine. A firewall needs to be recognized as a fortification. A honeypot is an ambush site or a delaying obstacle in cyberspace. Scanning is reconnaissance, and networks are areas of responsibility. Cybersecurity Service Providers (CSSP) and Enterprise Operations Centers are cyber defense battalions, brigades, or higher. Offensive cyber mission teams conduct raids, strike targets, and execute active defense missions using preemptive attacks. The Internet is no longer just the Internet; it's the battlefield.

Militarizing cyber symbols will give the cyber warrior insight into the parallel and analogous activities performed in other domains and will allow joint commanders to understand just what is happening in the cyber fight. The general might be unclear on what "Mimikatz" is or how it got through the firewall, but he will intuitively understand red arrows bypassing his fortifications and driving deep into his cyber key terrain. Commanders will soon learn to discern which cyber-related decisions are risky and which are not. The cyber battle, currently fought apart from the air-land-sea battle, must and will gradually be integrated into joint operations as doctrine evolves.

Doctrine will ultimately benefit from cyber symbols that conform to a joint standard. Cyber warriors already know the basic tactics for securing the battlefield, but an inability to visualize the battle hampers creation of a nuanced flow of cyber combat. At the opposite end of the spectrum, Joint Publication 3-12 (JP 3-12), *Cyberspace Operations*, brought some order to cyber command and control, but the paucity of operational doctrine has left a gulf between the tactical and the strategic. With proper symbols, concepts can be developed, presented, understood, and evolved by the joint community. Standards can be created, e.g., how many defenders are necessary for 50,000 accounts? Basic military precepts such as tempo and attrition can be addressed in a cyber context. Operational requirements can be identified, and the systems and equipment needed to meet that need can be acquired.

For cyberspace to truly become a warfighting domain, with all that entails, development of symbols that conform to joint standard is a necessary first step. To meet this need, IDA researchers developed a symbol set that is compliant with MIL-STD-2525, logically consistent, and capable of displaying the nuances of cyberwarfare to warfighters from all domains.

Terrain Graphics

JP 3-12 divides cyberspace into three layers: the physical, the logical, and the persona. The physical layer is the hardware, located in the physical domain, on which the other two layers exist. The physical layer is not cyberspace terrain itself. Symbols for physical equipment already exist in MIL-STD-2525D, *Joint Military Symbology*, and are not addressed here.

The logical layer is where cyber terrain exists, and the primary cyberspace terrain feature is the network, a collection of devices that implement applications, services, and data stores. Networks are the cyberspace equivalent to areas of operations in the physical domain. When protected by a firewall and monitored by intrusion-detection services at ingress points, a network becomes fortified and has a sensor line; when guarded by cybersecurity service providers and local cyber defenders, it is analogous to the most common command and control area designation: the operational area (OA).

We depict individual networks with a unique color-coded boundary line that represents the extent of the IP address space within it (see Figure 1).

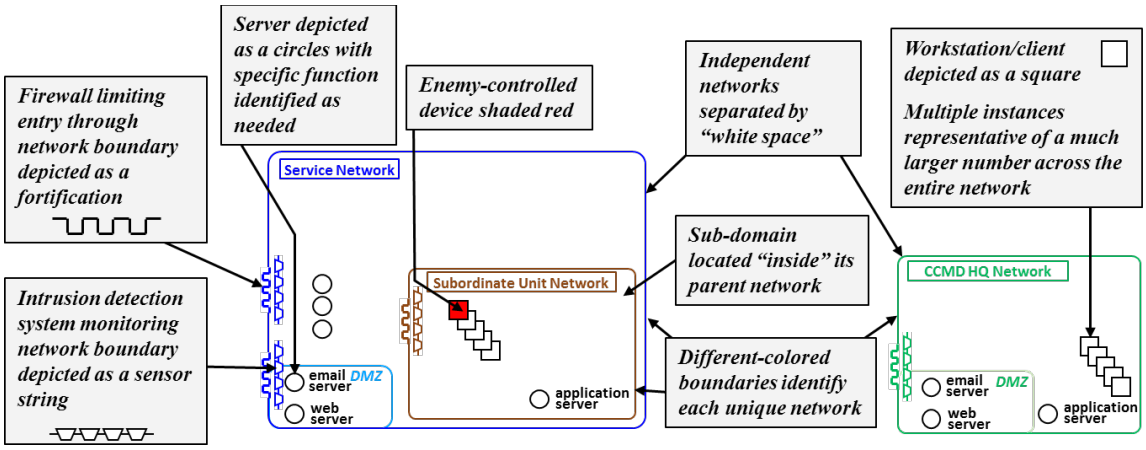

Figure 1. Cyberspace terrain description – networks and common features

For clarity, we typically depict only sufficient numbers of devices necessary to describe the planned or observed cyberspace operations, or to convey understanding of the nature of the terrain. For instance, if only one device out of hundreds on the network is attacked, we may choose to show that device alongside half a dozen others, often with a note that the small number of devices depicted are representative of many more.

Because of the nature of cyberspace, the distance between and the relative positioning of unique independent networks have little meaning in operational graphics depictions. However, the relationships between networks, such as one being a subdomain of another, *is* important, so we depict subdomains as existing completely within their parent networks.

Devices in cyberspace generally function simultaneously as terrain features upon which forces maneuver and as installations (which provide necessary supply, transportation, command and control, defensive, surveillance, or other warfighting functions); thus they have no clear analogies in the physical domain. We adopt common network diagram symbols in simplified form, depicting an individual workstation or client as a square and a server as a circle. However, we depict two specialized devices (and the functions they perform) that are nearly always present in cyber battles with unique symbols: the firewall is represented as fortification, and the intrusion detection equipment and services are represented as a string of sensors.

Red shading represents devices that have fallen under enemy control in some way. In some instances, red shading may be used to represent enemy control over an entire network.

Persona and Credential Graphics

The persona layer is the means by which personnel and units operate in cyberspace. JP 3-12 rightly asserts that the cyber-persona layer requires a higher level of abstraction but introduces confusion when it states that the persona layer consists of the people actually on the network. People do not exist in cyberspace, of course. Accounts and

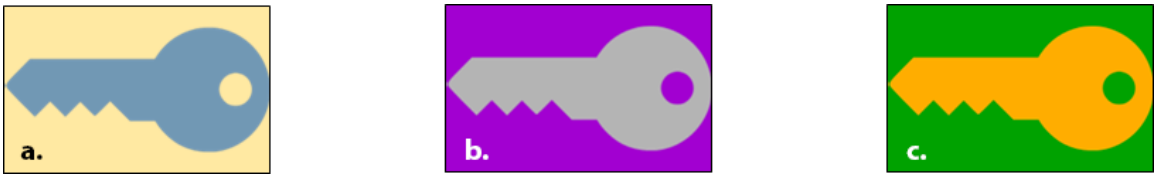

a. user-level credential with privileges in network identified by yellow boundary

b. system-level credential with privileges in network identified by purple boundary

c. domain-level credential with administrator privileges across network identified by green boundary

Figure 2. Notional cyber credential icons

their associated credentials (e.g., usernames, passwords, Common Access Card PIN) are the primary cyber entities that operators use to execute administrative actions, domain control, user activity, printer access, or any number of function-related activities—a network user account is a piece of cyber equipment that allows the operator to conduct email, use an Microsoft Office application, or communicate with other accounts. Similarly, in the air domain, a pilot (the operator) uses an F-22 (a piece of equipment) to conduct a variety of air superiority missions.

The difference is that the F-22 operator is physically paired with his equipment in the air domain itself whereas the cyber operator resides in the physical domain (where the physical layer of cyberspace exits) and conducts his mission in the cyberspace domain via the logical and persona layers. Cyber units thus have a foot in two domains, the living operators and physical layer hardware in one domain and the mixed types of accounts, credentials cyber actions, and missions in another.

Credentials are the "keys" to the cyber equipment and associated accesses and privileges. An adversary who gains credentialed access to a domain administration account is able to use the privileges associated with this account to control all of the key terrain—accounts, servers, data, and applications—in that OA. Different key symbols reinforce this point: blue for user-level, silver for system-level, and gold for domain-level privileges. A colored border around the key indicates the domain or network to which the privileges pertain (see Figure 2).

Unit Graphics

MIL-STD-2525D prescribes the use of specific frames for icon-based symbols to depict the identities of units operating in the land, sea, air, space, and subsurface physical domains. We adopt a regular hexagon frame to depict units in cyberspace (i.e., the logical and persona layers). We use standard colors for friendly and hostile entities and rotate the hexagons by 30 degrees to depict hostile units (Figure 3).

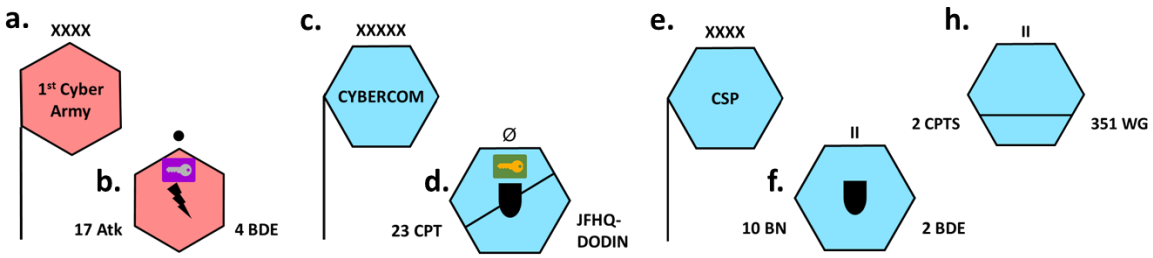

a. adversary HQ

b. adversary squad-level OCO unit with captured system admin credentials

c. U.S. Cyber Command HQ

d. friendly DCO unit with reconnaissance capabilities that has been granted domain admin credentials/authorities

e. friendly Cybersecurity Service Provider HQ

f. friendly DCO unit

h. friendly DODIN Ops cyber unit

Figure 3. Notional cyber unit icons

An icon, the innermost part of a symbol that provides an abstract pictorial or alphanumeric representation of units, equipment, installations, activities, or operations, must necessarily represent the unique nature of cyberspace units. Although cyber units may be equipped with specific "platforms" and trained for very specialized, unique missions at the lowest tactical levels, in general the diversity of the functions that cyber forces are capable of prohibits unique categorization by unit type based on specific equipment or mission as is typical in the physical domains (e.g., infantry versus mechanized infantry versus armor battalions, F-22 versus E-3 versus KC-135 squadrons). Instead, we use symbols that identify cyber units based on which of the three general mission categories from JP 3 12 they typically perform: Offensive Cyberspace Operations (OCO), Defensive Cyberspace Operations (DCO), or DoD Information Network (DoDIN) Operations. A lightning bolt identifies OCO units, a shield icon identifies DCO units, and existing support unit iconography identifies DoDIN Operations units. Cyber units performing the "detect" mission are depicted with a diagonal slash across the frame, similar to the use of a slash to denote "reconnaissance" capabilities in the physical domains.

We chose to adopt the existing echelon representation (used primarily in representing land force units) and apply it using the official designations of cyberspace units, with cyber protection teams representative of the lower echelons of friendly cyber forces typically portrayed, and U.S. Cyber Command as the top echelon.

Mission Graphics

In addition to the potential utility of adapting general offensive graphics (axis of advance, direction of attack), general defensive graphics (fortified line for firewall, sensor outpost for monitored intrusion detection device/system), and supply graphics (main supply routes or lines of communication for data flows), the traditional definitions of tactical mission graphics can be modified to depict actions in cyberspace. Potential adaptations of these graphics to cyberspace are provided in Table 1. The Doctrinal Description is as described and depicted in various DoD sources, including MIL-STD-2525D.

Other tactical tasks potentially useful for describing cyberspace actions were omitted from Table 1 for the sake of brevity or because no associated operational graphic exists: Control, Counter-reconnaissance (Area Security, Local Security), Disengage, Follow and Assume, Follow and Support, Defeat, and Suppress.

Table 1. Adaptation of Tactical Task Graphics to Cyberspace

Tactical Task	Operational Graphic	Doctrinal Description	Potential Use in Describing Cyberspace Operations
ACTIONS BY FRIENDLY FORCE			
Attack by fire		The use of direct fires, supported by indirect fires, to engage an enemy force without closing with the enemy to destroy, suppress, fix, or deceive that enemy.	Overt actions where an origination (or interim relay) point can be determined, such as Distributed Denial of Service attacks, broad intrusive scans, where these actions create the intended effect on the target.
Breach	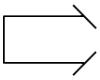	Break through or establish a passage through an enemy defense, obstacle, minefield, or fortification.	Non-credential-based access (e.g., penetration through a firewall, using an exploit or hacking tradecraft).
Bypass		Maneuver around an obstacle, position, or enemy force to maintain the momentum of the operation while deliberately avoiding combat with an enemy force.	Credential-based access (use captured credentials for login).
Clear		Remove all enemy forces and eliminate organized resistance within an assigned area.	Comprehensive scans and forensics, removing all malware and adversary points of presence and external connections.
Control	n/a	Maintain physical influence over a specified area to prevent its use by an enemy or to create conditions necessary for successful friendly operations.	Standard cybersecurity mission to protect a domain, typically assigned to a CSSP.
Counter-reconnaissance (Screen)	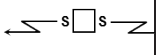	Provide early warning to the protected force.	Detection activities on a boundary or domain.

Table 1. Adaptation of Tactical Task Graphics to Cyberspace (continued)

Counter-reconnaissance (Guard)		Protect the main body by fighting to gain time while also observing and reporting information and preventing enemy ground observation of and direct fire against the main body. Units conducting a guard mission cannot operate independently because they rely upon fires and combat support assets of the main body.	Domain-wide detection and hunt-type activities by a Cyber Protection Team (CPT) or local defensive unit, augmenting the capabilities of a CSSP.
Counter-reconnaissance (Cover)	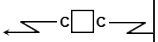	Protect the main body by fighting to gain time while also observing and reporting information and preventing enemy ground observation of and direct fire against the main body.	Domain-wide detection, hunt, and reposturing of defensive boundary controls by a CSSP.
Exfiltrate	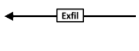 (No symbol exists. Symbol shows the flow of exfiltrated data, a substantial deviation from the existing definition of this task.)	Remove soldiers or units from areas under enemy control by stealth, deception, surprise, or clandestine means.	Movement of data from its original location to a location under enemy control, typically by means of stealth, deception, or clandestine means.
Occupy		Move a friendly force into an area so that it can control that area. Both the force's movement to and occupation of the area occur without enemy opposition.	Deployment of a CPT to a domain in advance of suspected adversary activity.
Retain		Ensure that a terrain feature controlled by a friendly force remains free of enemy occupation or use.	Defense of a network device or domain to prevent any adversary access.
Secure		Prevent a unit, facility, or geographical location from being damaged or destroyed as a result of enemy action.	Defense of a network device or domain to prevent an adversary from making any changes to data or functionality.
Seize		Take possession of a designated area by using overwhelming force.	Gain control of a device, network, data, or credentials. In cyberspace, two opposing forces *may* have simultaneous control of any or all of these assets.
Support by fire		A maneuver force moves to a position where it can engage the enemy by direct fire in support of another maneuvering force.	Overt actions where an origination (or interim relay) point can be determined, such as Distributed Denial of Service attacks and, broad intrusive scans, and where these actions are designed to set the conditions for success for the primary attack actions.

Table 1. Adaptation of Tactical Task Graphics to Cyberspace (continued)

EFFECTS ON ENEMY FORCE			
Block		Deny the enemy access to an area or prevent the enemy's advance in a direction or along an avenue of approach. Also an obstacle effect that integrates fire planning and obstacle efforts to stop an attacker along a specific avenue of approach or prevent the attacking force from passing through an engagement area.	Use or modification of blacklists, whitelists, access control lists, routing policies, credentials (username-password pairs or machine-issued), filters on firewalls, DNS servers, domain controllers, web servers, email servers, or others to prohibit or terminate access based on specific criteria.
Canalize	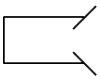	Restrict enemy movement to a narrow zone by exploiting terrain coupled with the use of obstacles, fires, or friendly maneuver.	Use of routing policies, honeypots/honeyports/ honeynets, or other defensive techniques to direct potential adversary traffic to desired network locations.
Contain	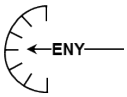	Stop, hold, or surround enemy forces or to cause them to center their activity on a given front and prevent them from withdrawing any part of their forces for use elsewhere.	Not strictly possible in cyberspace, since forces exist as a function of effort being expended. However, could be used to indicate quarantine of malware or emails.
Destroy	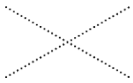	Physically render an enemy force combat-ineffective until it is reconstituted. Alternatively, to destroy a combat system is to damage it so badly that it cannot perform any function or be restored to a usable condition without being entirely rebuilt.	Deleting all files from a server, flashing BIOS or firmware, or causing physical damage to industrial control systems.
Disrupt	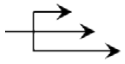	Integrates direct and indirect fires, terrain, and obstacles to upset an enemy's formation or tempo, interrupt the enemy's timetable, or cause enemy forces to commit prematurely or attack in a piecemeal fashion.	Interrupting connections periodically, enforcing time limits on sessions, or actions that require an enemy to repeat previous steps, upset an enemy's tempo, interrupt the enemy's timetable, or cause the enemy's efforts to proceed in a piecemeal fashion.

Table 1. Adaptation of Tactical Task Graphics to Cyberspace (continued)

EFFECTS ON ENEMY FORCE			
Fix		Prevent the enemy force from moving any part of that force from a specific location for a specific period.	Not strictly possible in cyberspace, since forces exist as a function of effort being expended, but used to indicate actions that require an enemy to focus effort to restore function (e.g., reboot a domain controller or data server following an induced system crash); to expend much greater effort than planned to obtain an objective (e.g., consuming attacker resources using a realistic honeynet); or to refrain from using capabilities for fear of detection (e.g., refrain from activating implants because of increased random scans for active malware).
Interdict		Prevent, disrupt, or delay the enemy's use of an area or route.	Denial of network (data transport) services, or limiting access to services.
Isolate		Requires a unit to seal off—both physically and psychologically—an enemy from sources of support, deny the enemy freedom of movement, and prevent the isolated enemy force from having contact with other enemy forces.	Removal of a device infected with malware from the network, moving a phishing email from the server to a forensics sandbox.
Neutralize	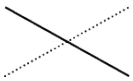	Render enemy personnel or materiel incapable of interfering with a particular operation.	Any action taken against another cyberspace *unit* that prevents it from using its offensive or defensive capabilities (e.g., interrupt the sensor feeds from a target domain to the responsible cyber defense unit).

Putting It All Together

These basic building blocks allow portrayal of cyber battles in a straightforward manner and present the action to the joint warfighter in a familiar format. The symbol set is still small—units, terrain, command and control, attack vectors—but capable of providing insights the commander needs for a rudimentary situational awareness of his OA. For example, battle maps with an attack arrow showing an enemy task force masquerading as friendlies and penetrating a fortification to pass undetected through sensors provide the joint force commander an enormous red flag that signals risk to the mission, which has been missing from the cyber portion of joint warfighting.

Figures 4, 5, and 6 depict the progression of a notional battle in cyberspace, from the initial assignment of defensive forces to their areas of responsibility, followed by the attacker's preparatory reconnaissance operations, and culminating in the penetration of defenses and the attacker occupying defended territory and postured to conduct follow-on operations. The astute reader will notice the similarities to historical depictions of Civil War battlefields, which motivated the development of these graphics to clearly depict complex, sequential actions over extended durations.

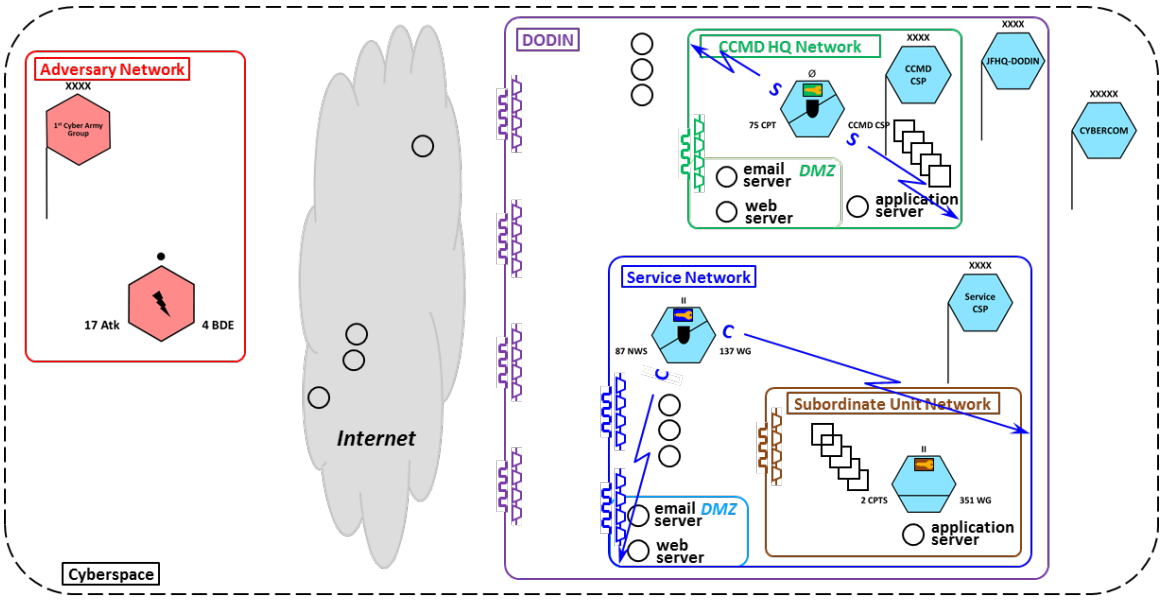

Figure 4. Notional cyberspace terrain showing boundaries, units, and defensive tasks

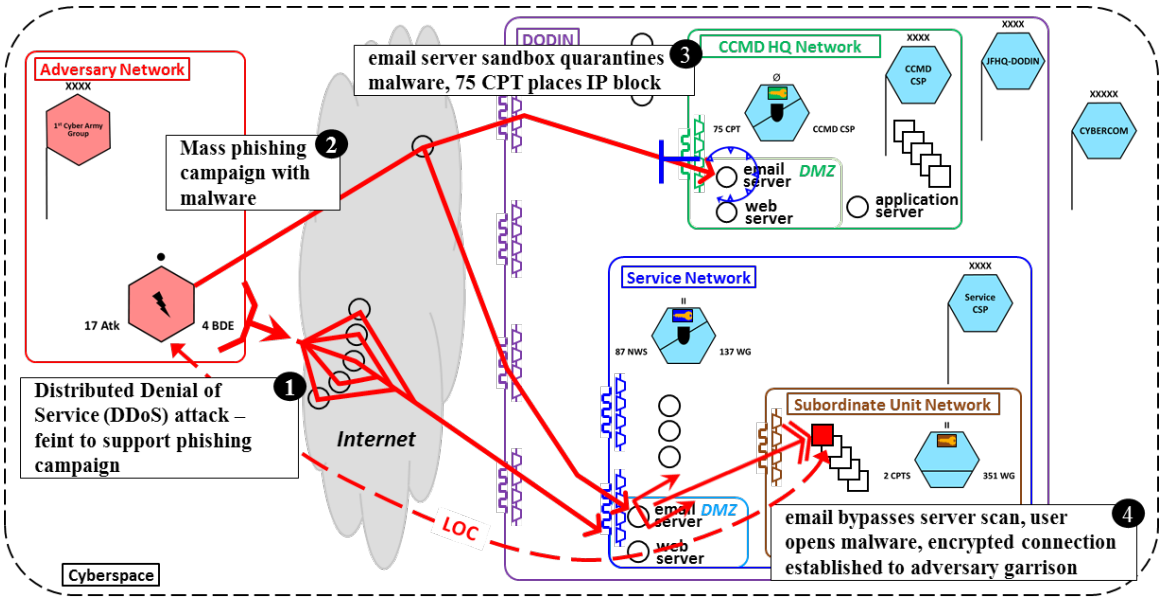

Figure 5. Sequential actions in the intial adversary assault: a feint, a blocked phishing attack, and a successful bypass of the defenses that gains control of friendly terrain

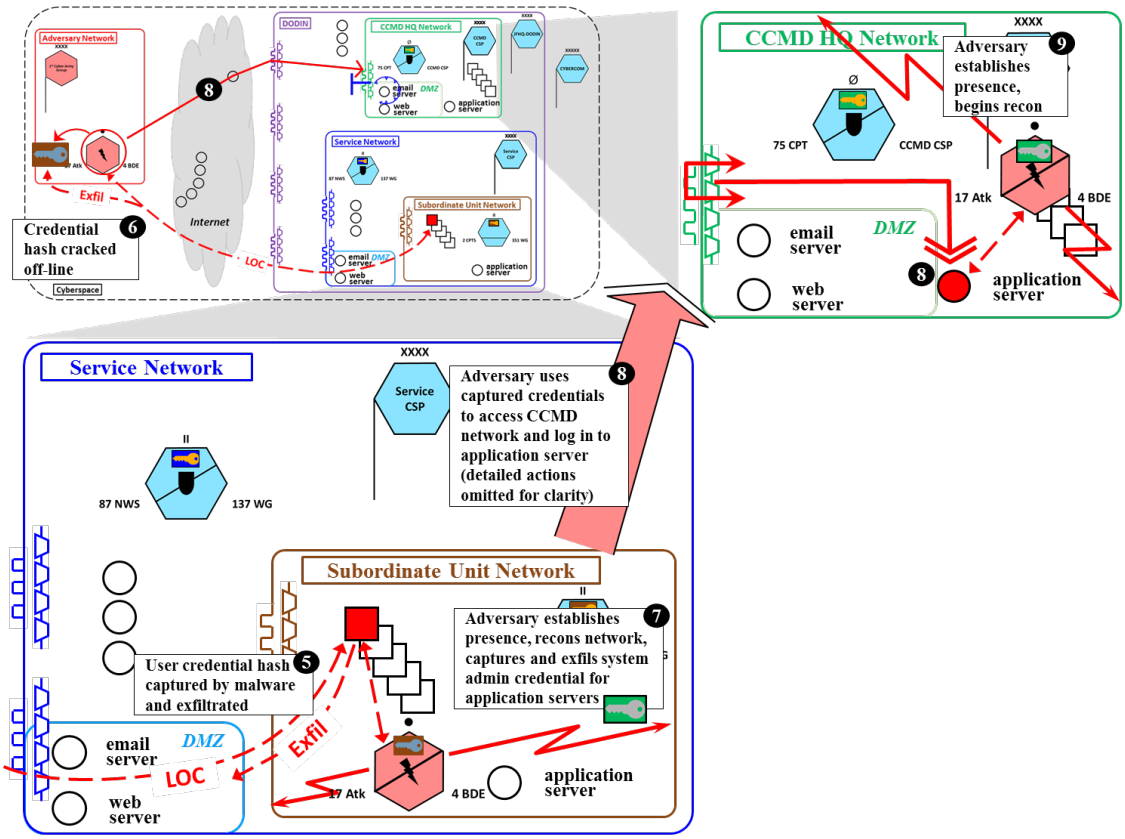

Figure 6. Subsequent adversary actions on friendly terrain: seizing of credentials, reconnaissance, and lateral movement within and between networks

Conclusion

Cyberspace operational graphics will allow cyber planners and operators to convey mission-relevant information to warfighters who are unfamiliar with the technical details of cyberspace. Military tasks, missions, and operations share commonalities regardless of the domain in which they take place, and leveraging warfighter familiarity with the common language that has evolved to describe them will enhance rapid understanding and decision-making.

Using operational graphics to describe cyberspace actions should lead to the identification of parallels and analogies in the physical domains that could potentially be implemented in cyberspace operational doctrine. For instance, the doctrinal concepts of culmination and attrition that are critical to operational campaign design and execution in the physical domains may finally be examined fully for application in the cyber domain. Ultimately,

the joint commander will have at his disposal a coherent body of operational doctrine and the accompanying graphics that enable him to understand, plan, and fight the cyber battle.

IDA hosted the first-ever DoD Cyber Symbology Workshop in February 2019, and the symbol set is in the process of being refined and formally incorporated into MIL-STD-2525 by DoD.

Reference

MIL-STD-2525D, Joint Military Symbology, 10 June 2014; FM-102/MCRP 5-12A, Operational Terms and Graphics, 2 February 2010 (incorporating Change 1); FM 3-90-1, Offense and Defense, Volume 1, March 2013; FM 3-90-2, Reconnaissance, Security and Tactical Enabling Tasks, Volume 2, March 2013.

Mr. Erick D. McCroskey is an Adjunct Research Staff Member in IDA's Operational Evaluation Division. He holds a Master of Science in applied engineering (physics) from the University of California, Davis, and a Master of Arts in military arts and sciences from the School of Advanced Military Studies, Fort Leavenworth.

Mr. Charles A. Mock is a Research Staff Member in IDA's Operational Evaluation Division. He holds a Master of Science in computer science from the Naval Postgraduate School.

Operationally Assessing Cyber Defenses

Allison Goodman

The Challenge: When Congress directed the Director, Operational Test and Evaluation to establish a Cybersecurity Assessment Program, DOT&E asked IDA to help plan realistic data-based cyber evaluation events during operational exercises, execute the events, and provide rigorous post-event analyses. Analytical methods for these evaluations must continue to evolve and become more rigorous as our cyber adversaries become more sophisticated.

DOT&E conducts data-based cyber evaluations during Combatant Command and Service-level exercises to... determine the effect of cyber activities on operational missions.

IDA supports the Director, Operational Test and Evaluation (DOT&E) on the congressionally mandated Cybersecurity Assessment Program. Through this program, DOT&E conducts data-based cyber evaluations during Combatant Command and Service exercises to understand how a cyber adversary can attack and compromise networks, characterize the defensive response, and determine the effect of cyber activities on operational missions. The program uses DoD Cyber Red Teams to portray a live, thinking, cyber adversary, who join with the traditional exercise Opposing Force to target critical Combatant Command Missions within the context of exercise scenarios designed to exercise operational plans. IDA analyses in support of DOT&E have produced recommendations for both local and department-wide defensive approaches and vulnerability mitigation strategies. This article describes how IDA supports the DOT&E Cybersecurity Assessment Program and the analytical methods used to conduct these data-based evaluations.

IDA support spans the lifecycle of the exercise: planning, execution, and post-assessment reporting. During the planning events, IDA researchers help scope the cyber component of the event while still ensuring that the cybersecurity assessment will not negatively affect the exercise training objectives.

During the exercise, the operational test agencies collect data on four main cybersecurity functional areas: Protect, Detect, Respond, and Recover (PDRR). Data collection focuses on both those executing the exercise mission (operators and cyber defenders) and the opposing force portrayal (Red Team). IDA researchers are on site during exercise execution, ensuring data accuracy and completeness, as well as maintaining situational awareness for the post-assessment analysis and reporting.

The collection of PDRR data in the context of the exercise allows IDA researchers the ability to focus on attack threads, defensive responses, and mission effects. Attack threads detail each step in an attack, starting from intrusion and ending at either mission effect or detection. Figure 1 shows the intended outcomes of two notional cyber-attacks, which will be used to illustrate the analytical process.

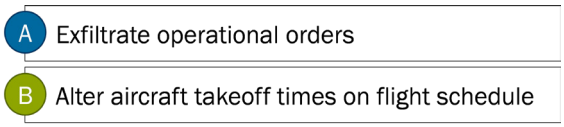

Figure 1. Notional Cyber Attack Thread Outcomes

Red Teams provide detailed information about each action taken during the exercise, including methods and tools used. IDA researchers organize these actions by those leading to the identified attack thread outcomes and map the progression of the cyber-attack from ingress to conclusion. Figure 2 shows the simple and notional Attack Thread A to illustrate this mapping. In this example, the Red Team affects the confidentiality of the operational orders by exfiltrating them from the system and network. Defenders did not detect the Red Team movement through the network or the exfiltration of data, and therefore do not appear in the notional thread.

Attack Thread B illustrates the combination of the Red Team actions with the cyber defender actions (Figure 3). In this example, the end user detected and reported the modified takeoff times, and the cyber defenders responded by identifying and blocking the originating IP address. Data for these cyber actions, detections, and response comes from multiple sources, which IDA researchers combine to present the end-to-end picture of each cyber-attack.

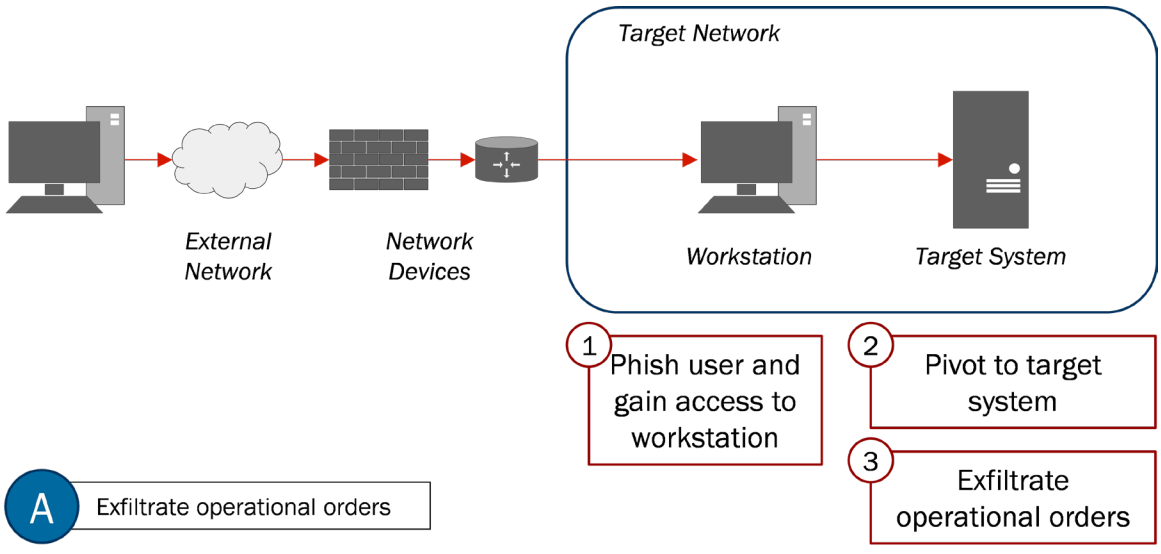

Figure 2. Notional Attack Thread A

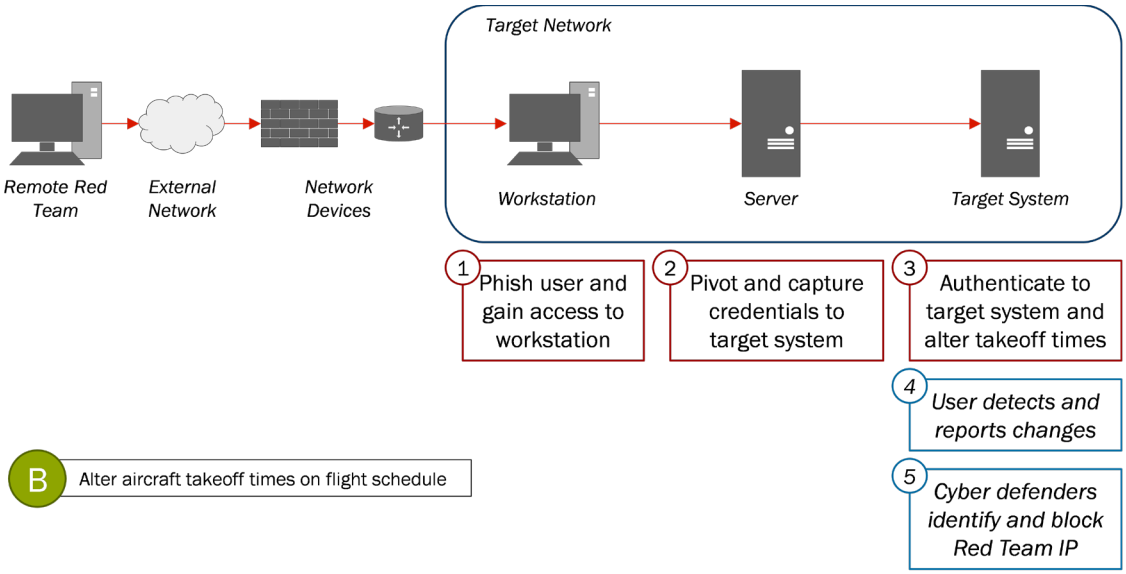

Figure 3. Notional Attack Thread B, including cyber defender actions

Next, IDA researchers incorporate the effect on the operational mission from the exercise scenario, providing context to the outcomes of each attack, as applicable. Following Attack Thread A, if the Red Team exfiltrates the operational orders after they were already executed, this has little effect on the overall mission. However, if the Red Team exfiltrates them prior to execution of the orders, the opposing force has knowledge of future friendly force activity, providing the opportunity to disrupt operations.

To provide further context, IDA researchers also determine the capability level required to execute each attack thread by evaluating the knowledge; tools, techniques, and procedures (TTPs); and planning required to execute each attack thread. The capability is rated on a four-level scale ranging from nascent to advanced for each of these categories and their sub-categories. Table 1 shows the criteria for each capability rating by categories. The circles indicate the notional capability breakdown for Attack Threads A and B. The level of capability required to achieve a particular attack thread is then the greatest capability level required across all categories. Therefore, Attack Thread A required Limited capability to achieve and Attack Thread B required Moderate capability to achieve.

These analyses provide Combatant Commands and the Services with not only an analysis of network vulnerabilities, but also the potential effects that vulnerabilities could have on their missions and the capabilities required to achieve those effects.

Table 1. Notional Capability Required to Complete Attack Threads A and B

		Nascent	Limited	Moderate	Advanced
Knowledge	General Systems	Common OS (Windows client, Linux), and software applications (Adobe, Oracle), consumer-market hardware (PCs, home routers), common network and data protocols (IP, Ethernet, 802.11), general-purpose languages (Python, Java, SQL), common OS-specific languages (Unix shell, PowerShell), public cryptography and standard authentication (PGP, NTLMv2, Kerberos) B	Commercial enterprise OS (MS Server, virtualization environments), industry market network OS (Cisco IOS, Juniper) and devices (routers, proxies, VPN), defensive devices (IDS, firewalls), cellular data protocols (GSM, 4G LTE), common firmware (BIOS), common architecture assemblers (Intel, ARM, MIPS), token-based authentication (CAC, ActiveID) A	Common military software (GCCS, HBBS, TBMCS), less-common network and data protocols (tactical data links, radio, CAN bus, other MIL-STD interfaces), embedded systems (PLCs, digital signal processors) and software (embedded C, RTOS), specialized firmware (fuzes, avionics), server/ military assemblers (SPARC, MIL-STD 1750A), biometric-based authentication	Restricted and highly classified military systems, software, and weapons platforms, classified cryptography (NSA Type 1) and associated hardware (TACLANE), cross-domain devices (Radiant Mercury, ISSE Guard)
	Target Network and Systems	Information about target environment found from commonly available open sources (commercial Internet, literature) or from external reconnaissance of target network and systems	Knowledge of network and system specifications (individual user account information, hostnames, IP address of few systems) and type/configuration of host-based defenses equivalent to an authorized user in the target environment A B	Knowledge of network and system specifications (configuration settings, software inventories) and type/configuration of networked defenses (IDS, ACLs) equivalent to an authorized Administrator in the target environment	Knowledge of network and system specifications (network architecture, Domain-wide configurations and user account information) and defenses (full defense in depth) equivalent to an authorized Domain Administrator in the target environment
	Target Operations	Information found from commonly available open sources or from external reconnaissance of target organization	Knowledge from more specialized literature or equivalent to prior experience with target operations, including key information or supporting systems A	Knowledge equivalent to substantial prior experience with target operations, including work flow and sub-task objectives B	Knowledge of current target operations equivalent to an experienced authorized operator
Tools	Software and Hardware	Freeware (Kali, Scapy, Poison Ivy) and inexpensive commercial tools (Retina, Cobalt Strike), public exploits of known vulnerabilities (Metasploit, w3af), inexpensive hardware (PCs, Yellowjacket, rogue WAPs like PWN Plugs, physical access tools, connectors) A	Commercial software (Core Impact, Metasploit Pro), 0-day exploits of less common/more vulnerable software (Adobe, MAC OSX), custom software (kernel rootkits, C2 agents) and hardware (GPU clusters, covert rogue WAPs) costing $10,000s or dozens of man-hours B	0-day exploits of more common/less vulnerable software (Windows, iOS), custom software (polymorphic malware, covert remote access tools and loggers, boot sector/firmware rootkits, forged SSL certificates) and hardware (rogue MIL-STD WAPs) costing $100,000s or hundreds of man-hours	0-day exploits of restricted military systems and industrial control systems, custom software (firmware-resident malware, high-level programming languages) and custom hardware (covert RF WAPs, chipset backdoors, TEMPEST devices), costing $1,000,000s or thousands of man-hours
Operations	TTPs	No demonstrated stealth, non-attribution or efficient use of resources A B	Low degree of stealth (C2 over uncommon protocols, changing signatures or running tools in memory to avoid common A/V, rootkits), non-attribution (log purging, IP/MAC spoofing, TOR), or efficiency in use of resources consistent with intent	Some degree of stealth (C2 with custom encoding, disabling A/V or IDS), non-attribution (code obfuscation, fast-fluxing), or efficiency in use of resources consistent with intent	High degree of stealth (strategic onetime use C2, full control of defensive infrastructure), non-attribution (false flag operations), or efficiency in use of resources consistent with intent
	Planning	Opportunistic actions, no planning	Intent and short-range plans formed on-the- fly as needed A B	Organizes (one or more) operations with specific target systems and associated effects on target organization	Organizes multiple operations against separate targets, synchronizing timing, accesses, and planned second-order effects

IDA continually evolves this methodology as the cyber-attacks and defenses grow in complexity. IDA researchers identify data gaps and ensure that upcoming assessments fill those gaps. Each exercise presents the opportunity to research new questions and provide more insight into the state of cybersecurity across DoD.

Dr. Allison Goodman is a Research Staff Member in IDA's Operational Evaluation Division. She holds a Doctor of Philosophy in biomedical engineering from Virginia Polytechnic Institute and State University.

Informing a Defensive Strategy by Analyzing Reactive Cyber Intrusion Detection

V.V. Kulkarni and S.C. Whetstone

The Challenge: The Department of Defense Information Network (DoDIN) is under constant attack from adversaries ranging from independent hackers to sophisticated nation-states. Breaches in computer networks can occur through phishing campaigns, exploiting vulnerabilities via hacking tradecraft, malicious insiders, and surreptitious physical access to network infrastructure.

Breaches in computer networks can occur through phishing campaigns, exploiting vulnerabilities via hacking tradecraft, malicious insiders, and surreptitious physical access to network infrastructure.

The Department of Defense Information Network (DoDIN) contains more than 4 million computers and 3 million users. The DoDIN has access points around the world, including on military bases, ships, aircraft, and cellular devices. Many sub-networks in the DoDIN transfer data over the same physical channels as the Internet, including fiber optic cables and satellite. With so many entry points, the DoDIN is under constant attack from adversaries ranging from independent hackers to sophisticated nation-states. Breaches in computer networks can occur through phishing campaigns, exploiting vulnerabilities via hacking tradecraft, malicious insiders, and surreptitious physical access to network infrastructure. Network administrators and defenders traditionally focus on protective defense, i.e., preventing initial network compromise. Unfortunately, persistent adversaries will inevitably find ways to breach protective defenses. It is therefore crucial that network defenders receive training on reactive defense: the ability to detect, respond to, and restore networks after initial compromise. This article describes how IDA's analytical framework for operational cybersecurity assessments of Combatant Command training exercises conducted between fiscal years 2014 through 2016 informed a defensive strategy.

The data collected during these assessments provide insights on both the attacker and defender actions. The attackers act as part of the assessment team and provide information regarding the individual actions, the linkages between actions, and their perception of success. The attack thread is the instrument used to probe and measure the detection capability of the cyber defenses. The defenses and defenders are under evaluation, and the data collection captures detection of the attacks, defensive responses, and operational effects.

Analytical Focus—Cyber Attack Threads

An attack thread is *a series of steps taken by an adversary encompassing the intrusion and subsequent exploitation of a network*. Attack threads end either with the adversary conducting an information effect on an objective node or with the network defenders stopping the adversary before they reach an objective node.

What factors are important in detecting an attack? Experience with assessments suggests two factors: type of access and type of tool used.

Detection Factor—Logical Access

Logical access to software is either authenticated or unauthenticated. Authenticated access involves presenting a credential, which the software checks and validates before granting access. Credentials come in many forms, such as usernames and passwords, or tokens and PINs. Unauthenticated access involves accessing software without presenting credentials. Adversaries gain unauthenticated access by techniques such as SQL injection, malicious file uploads, booting a workstation from an unauthorized DVD, buffer overflows, and other malformed requests.

Detection Factor—Tool Type

Cyber attackers use tools to perform actions. These tools are either native or foreign. A native tool is one that the network owners authorize for use on the network. Since many operating systems in the DoDIN run Microsoft Windows, many of Microsoft's command line tools and software packages are permitted and are considered native tools (Powershell for instance). Foreign tools are tools that network owners have not authorized for use on the network. Foreign tools include scanners, malware, viruses, beacons, and command and control software.

Cybersecurity principles suggest how detection will vary with these two factors. Defenses typically attempt to identify suspicious or unusual actions and alert defenders to investigate them. Authenticated access and use of native tools by definition are normal actions and thus typically generate no alerts, making attack actions with these characteristics harder to detect, as illustrated in Figure 1.

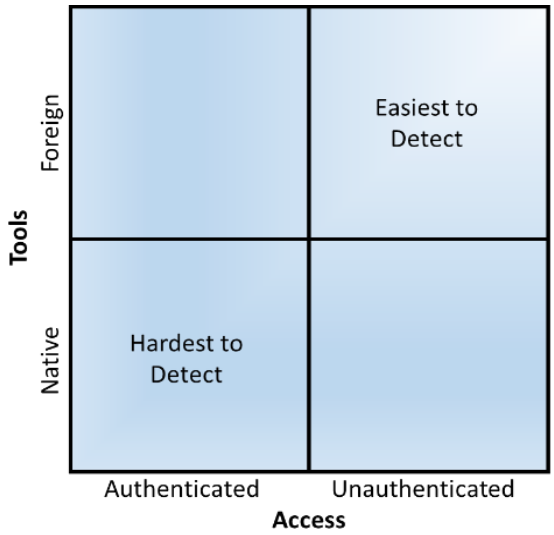

Figure 1. Network defenders have a higher likelihood of detecting the Red Team when they use unauthenticated access and foreign tools in their attack threads

This simple model suggests a two-fold defensive strategy for improving detection. First, force the adversary to operate in the portion of the space where detection is easier by using more foreign tools and unauthenticated access.

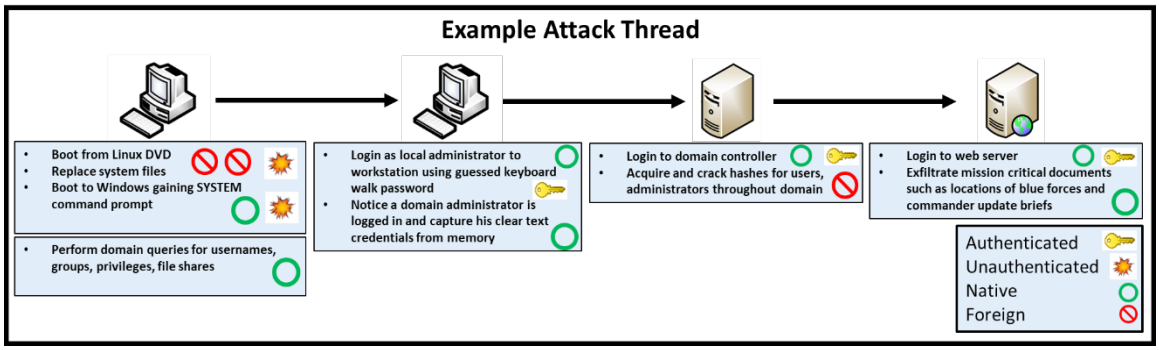

Figure 2. An example Red Team attack thread

Second, structure the defenses to reduce the portion of the space where actions are difficult to detect by improving detection of native tools and authenticated accesses by unauthorized users or adversaries.

The operational assessments use DoD Cyber Red Team attacks to stimulate and gather data about the defenses. Figure 2 shows an example attack thread with the Red Team starting from physical access to a workstation (without possessing login credentials), escalating their privileges to domain administrator, and subsequently stealing mission-critical documents from a file server undetected.

Strategy Implementation—Attack Thread Characterization Metrics

The details of the attack threads provide insight on how well the network defenses are forcing an adversary to rely on detectable actions during cyber-attacks. The ratio of actions within an attack that used unauthenticated accesses and foreign tools to the total number of actions in that thread captures the ability of the attacker to operate where detection is less likely.

$$metric\ 1 = \frac{\#\ of\ unauthenticated\ accesses}{total\ \#\ of\ accesses}$$

$$metric\ 2 = \frac{\#\ of\ actions\ using\ a\ foreign\ tool}{total\ \#\ of\ actions}$$

The approach is to compute the metric for each observed attack thread. For example, the attack thread depicted in Figure 2 contains five accesses, two of which are unauthenticated, yielding a score of 2/5 or 40%. Of 10 total actions, 3 use foreign tools, yielding a score of 3/10 or 30%. These metrics provide insight on the success in forcing attackers to operate where defenses can detect them. The metrics enable a comparative analysis to determine how changes in the network, perhaps over time or when changing a detection device, affect the attacker actions.

Next consider the ability to measure the defenses themselves.

Defensive Performance—Logistic Regression

A logistic regression predicts the probability to detect future attacks based on the ratios of unauthenticated accesses and use of foreign tools. This analytical approach uses a binary response: the defenders either

detected or did not detect the attack. Combining this binary response with the two metrics previously defined as the continuous variables in a logistic regression allows us to model the conditional probability that the defenders will detect an attack as a function of the fraction of unauthenticated accesses and foreign tool use. The conditional probability of detecting an attack thread given factors x and y is:

$$P(Detect|x,y) = \frac{e^{f(x,y)}}{1 + e^{f(x,y)}}$$

$$f(x,y) = \beta_0 + \beta_1 x + \beta_2 y + \beta_3 (x - \bar{x})(y - \bar{y})$$

In this model, the variables x and y represent metric 1 (unauthenticated access) and metric 2 (foreign tool use). β_0, β_1, β_2, β_3 are constants. $\bar{x}$ and $\bar{y}$ represent the averages of the two metrics over all attack threads. Each attack thread is represented by three values (x, y, Detected? (yes or no)), and statistical software determines the four constants in the logistic regression model.

This model contains only two factors: the percentage of unauthenticated accesses in an attack thread and the percentage of foreign tools used. Other factors affect detection rates as well. Network defenders may receive tips from intelligence reports for example. Additionally, network defenders may know the Red Team's IP address space prior to an assessment, skewing detection rates. Nonetheless, the two factors provide useful insight into observed defensive performance and how to safeguard an information network.

As an empirical model, the regression requires a data set of diverse attack threads. The analyst can then construct a two-dimensional space with contour lines showing how the probability to detect an attack varies with the characteristics of the attack, as illustrated in Figure 3. The operational assessments of training exercises in fiscal years 2014 through 2016 provided a data set to apply the methodology and provide insights to the DoD network defenses.

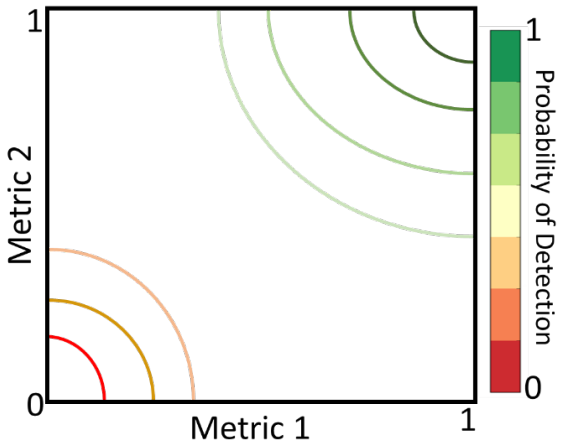

Figure 3. Conditional Probability to Detect defines an operational space with contours identifying regions of greater or lesser chance of detection. (Qualitative sketch only)

Application of the Framework

Logistic regression analysis of attack threads during training exercises confirmed insights on the detection performance of the observed cyber defenses. Although the specific performance values for the DoD networks are classified, observations regarding applicability of the framework are not.

The probability of detection varied with lower probabilities to detect in the lower left quadrant of the operational

space and higher probabilities in the upper right quadrant. This observation is consistent with expectations of difficulty in detecting actions with different types of access and tools. The logistic regression supports the two-fold defensive strategy for improving detection by denying an adversary the ability to operate where defenses are weak, and it identifies specific areas where defenses are less likely to detect an adversary.

The framework analytically confirmed the anecdotal observation that the network defenses are improving, but not enough to stop the cyber Red Teams. The logistic regression quantitatively confirmed an improving probability of detection over the period of the assessments, and that the Red Teams remained successful by adjusting their tactics to operate in the region of the operational space where the defenses were less likely to detect them.

In addition, the analytical framework can provide insights on how network design affects detectability. For example, assume that a network includes publically accessible websites that do not require authentication but store sensitive information. The network defenses would not necessarily alert on such accesses or detect attacks against those assets. The logistic regression should show a low probability of detection for attacks having a high fraction of unauthenticated access, which contradicts expectations from the simple model for difficulty in detection. Defenders with such knowledge could adjust their network design to minimize such publicly accessible websites or alter operational procedures to specifically monitor for such attacks to improve the probability of detection.

Integrating over the conditional probability to detect yields the total probability of detecting an attack:

$$P(Detect) = \iint P(Detect \mid x, y)\, p(x, y)\, dx\, dy$$

where $p(x, y)$ is the probability for an adversary to operate at point (x, y). In other words, $p(x, y)$ is the probability of a given attack thread to have metric 1 = x and metric 2 = y. This probability distribution is strongly dependent on adversarial tools, techniques, and procedures.

A simple approach is to assume the adversary is equally likely to operate anywhere in the operational space of Figure 3. In this scenario, $p(x, y) = 1$. The total probability to detect an attack can be calculated for different sets of attack threads to compare across time. Doing so shows that in fiscal years 2014 through 2016, the probability of a defender detecting an attack has risen. An analyst also could develop a tailored profile $p(x, y)$ from intelligence data for specific adversaries to estimate the expected defensive performance.

The analytical framework also leads to a general set of principles and recommendations for improving detection of attacks.

Force the adversary to use more foreign tools and unauthenticated access

Cyber adversaries seek valid credentials in order to blend in as authorized users. They may obtain credentials by cracking hashes

stored on disk, extracting clear-text credentials from memory, locating clear-text password files, guessing, and keylogging. Cracking hashes is trivial when users have weak passwords such as keyboard walks. Network administrators must routinely perform password audits to ensure that users have strong passwords.

Additionally, network administrators must remove all or encrypt all clear-text password files found on workstations and servers. If adversaries cannot acquire credentials, they must try unauthenticated access and are therefore easier to detect. Network administrators should also restrict the use of certain native tools. Publicized breaches of organizations show that adversaries commonly use Windows native tools such as Powershell, Procdump and PsExec. Although network administrators use these tools as well, normal users will not. Therefore network administrators should restrict the use of these tools.

Increase the detectability of native tools and credential misuse

Network defenders can configure host-based security system rulesets to flag upon any executable. They should therefore be aware of users using native tools that are not necessary in their day-to-day routines. For example, hackers execute Procdump on the Windows background running process lsass.exe in order to access clear text credentials from running memory. An innocent user would not do such a thing, and therefore network defenders can spot anomalous behavior performed by native tools. Furthermore, network defenders should be aware of which users routinely use remote login tools such as PsExec and which do not. Anomalous use of PsExec could be an indication of an adversary attempting to move laterally in the network with valid credentials.

Dr. Vikram V. Kulkarni is a Research Staff Member in IDA's Operational Evaluation Division. He holds a Doctor of Philosophy in physics from Rice University.

Dr. Shawn C. Whetstone is a Research Staff Member in IDA's Operational Evaluation Division. He holds a Doctor of Philosophy in nuclear engineering from the University of Michigan.

Cyberspace and Agility: Lessons from the Office of Personnel Management Breaches

David Alberts

The Challenge: When Office of Personnel and Management (OPM) data breaches affected 21 million Federal employees with security clearances, the Federal government turned to the Department of Defense to lead the near-term response and develop a long-term solution. The long-term solution requires a top-to-bottom transformation of the security clearance vetting ecosystem.

Not only were DoD employees and contractors the victims of these attacks, but the compromise of sensitive information collected during security background investigations has serious national security implications.

The Breaches and Their Significance

In June 2015, many current and former members of the Department of Defense workforce learned from two news reports that the Office of Personnel Management (OPM) had suffered at least two data breaches that could adversely affect them and their families. Federal officials characterized these breaches as among the largest breaches of government data in the history of the United States. As it turned out, the data compromised contained not only personally identifiable information (PII) such as Social Security numbers, but also fingerprint and security clearance-related information that employees, contractors, and applicants had provided on the Standard Form 86, Questionnaire for National Security Positions.

For the workforce, the breaches, subsequent announcements by Federal officials, Congressional testimony, and news stories regarding the breaches certainly did not instill a sense of confidence in the security of their information nor in the adequacy of the Federal response. Perhaps most alarming was that initial announcements of the number of people adversely affected grew from 4.2 million to more than 21 million.

Upon learning the extent of these attacks, senior leaders across the government dedicated themselves to better understanding the nature of these cyberattacks, taking steps necessary to prevent future intrusions, notifying the individuals affected, and offering them both identity protection and credit monitoring services. With more than 80 percent of the individuals who undergo security-related vetting, DoD was the agency most affected by the OPM breaches. Not only were their employees and contractors the victims of these attacks, but the compromise of sensitive information collected during security

background investigations has serious national security implications. One can easily imagine how this information could be used by a sophisticated nation-state actor.

The DoD Response

With the charge to respond rapidly, DoD was given the responsibility to lead an interagency effort to notify affected individuals and arrange for appropriate services. In July 2015, the DoD Chief Information Officer (CIO) formed a Notifications Tiger Team, led by the Deputy CIO for Cybersecurity, to plan and manage this effort. The objective was to notify those affected as promptly as possible, while protecting against any additional compromise of this information and the possibility of adversary counterintelligence exploitation.

IDA supported the DoD Notifications Tiger Team in several ways. These included identifying what tasks were needed and assessing, in real time, progress in identifying critical path items that needed immediate attention. IDA also provided support to the efforts to bring on board a contractor to provide credit monitoring and identity theft protections by developing a statement of work, formulating evaluation criteria, and serving as advisors to source selection and later security reviews to help ensure that personal and sensitive data would be appropriately protected.

The DoD-led interagency effort to notify those affected by the second OPM breach resulted in letters being sent to more than 90 percent of these individuals by mid-December 2015, less than six months after the Tiger Team was formed and less than three months from award of the credit monitoring and identity theft contract. The team met its principal objectives of notifying affected individuals while safeguarding sensitive information, despite an extremely aggressive schedule and the need to overcome a number of potentially mission-threatening challenges.

Need for Transformation

Based on what was learned from the post-breach efforts, it became clear that a top-to-bottom transformation of security clearance vetting ecosystem would offer the best chance to reduce the frequency and severity of future intrusions and compromises, as well as address significant shortcomings of the existing process and systems. IDA developed the following vision statement for a transformed vetting ecosystem:

> *A transformed end-to-end process supporting security, suitability and credentialing (SSC), Insider Threat, and CI that leverages the power of information technology, ubiquitous data, and automation operating in a secure, defended, agile, shared infrastructure.*

This vision was subsequently adopted by the interagency group that oversees and coordinates the Federal vetting enterprise.

As shown in Table 1, IDA compared and contrasted the capabilities and characteristics of current concepts based on legacy technologies to a transformed concept enabled by newer technologies.

This description of a transformed ecosystem vision involves more than a better information technology system and more than a re-engineered process. It is, first and foremost, an *agile* ecosystem that addresses the vetting challenge in a holistic manner, with the capability to learn from streams of real-time information, learn what works and what does not, and evolve its governance and design to seize upon opportunities for improvement and respond to stresses as circumstances change.

IDA continues to support the Under Secretary of Defense for Intelligence and the DoD CIO as DoD develops the new National Background Investigation Service. Also, OMB asked IDA to perform an agility-based analysis of the security clearance vetting ecosystem to ensure that agility is built into transformational efforts.

Table 1. Current Concepts vs. Transformed Concepts

Characteristics and Capabilities	Current Concepts and Legacy Technology	Transformed Concepts Enabled by Technology
Cybersecurity	patchwork	designed and built-in end-to-end defenses evolving with threats
Workflow	fixed periodic	dynamically reconfigurable anomaly triggers
Routine Tasks Information Gathering	many prescribed, manpower-intensive	fewer, tailored automated to the extent practical
Missions	one	multiple
Protection Levels	one	as many as needed
Ability to Change / Evolve	limited high cost – not timely	agile with evidence-based evolution

Dr. David Alberts is a Senior Fellow in IDA's Strategy, Forces and Resources Division. He holds a Doctor of Philosophy in operations research from the University of Pennsylvania.

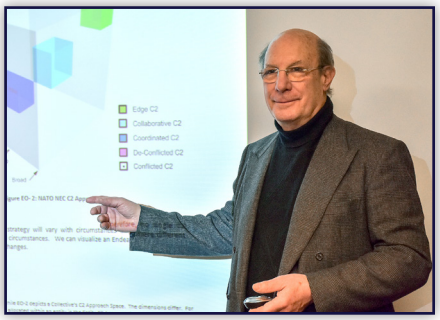

Past Issues

Challenges in Cyberspace: the Human Dimension
- Building the Cyber Warfare Force
- DoD's Cyber Workforce Challenges
- Staffing Cyberspace Operations
- Identifying Enlisted Recruits with the Right Stuff to Perform Cyberspace Operations
- Air National Guard Cyber Force
- Modernizing Air Force Cybersecurity Test and Evaluation

IDA Text Analytics
- IDATA Overview
- Data Exploration and Management of Defense Finance and Accounting Services Artifacts
- Extracting Structured Numerical Data from Large Quantities of Memoranda
- Implementing the Federal Advisory Committee Act
- Finding and Categorizing Recurring Reports to Congress
- Comparing the House and Senate Versions of the National Defense Authorization Act
- Discovering, Analyzing, and Understanding Improvised Explosive Device Documents
- Use of IDATA Capabilities for Social Media Analytics

Challenges in Cyberspace
- Cyberspace – The Fifth and Dominant Operational Domain
- Transitioning to Secure Web-Based Standards
- Information Assurance Assessments for Fielded Systems During Combat Command Exercises
- Supplier-Supply Chain Risk Management
- Internet-Derived Targeting: Trends and Technology Forecasting
- Training the DoD Cybersecurity Workforce

Multidisciplinary Research for Securing the Homeland – IDA and DHS: Beyond 15
- Countering Terrorism One Technology at a Time
- Does Imposing Consequences Deter Attempted Illegal Entry into the United States?
- Improving Shared Understanding of National Security and Emergency Preparedness Communications
- Foreign Counter-Unmanned Aerial Systems: Developments in the International Arms Market
- Operationalizing Cyber Security Risk Assessments for the Dams Sector
- Understanding the Juvenile Migrant Surge from Central America
- Implementing a Roadmap for Critical Infrastructure Security and Resilience
- Baselining: Application of Qualitative Methodology for Quantitative Assessment of Emergency Management Capabilities
- Analysis, Analysis Practices, and Implications for Modeling and Simulation
- Test and Evaluation for Reliability

Acquisition, Part 1: Starting Viable Programs
- Defining Acquisition Trade Space Through "DERIVE"
- Supporting Acquisition Decisions in Air Mobility
- Assessing Reliability with Limited Flight Testing
- Promise and Limitations of Software Defined Radios
- Implications of Contractor Working Capital on Contract Pricing and Financing
- The Mechanisms and Value of Competition
- Early Management of Acquisition Programs

Acquisition, Part 2: Executing and Managing Programs
- Cost Growth, Acquisition Policy, and Budget Climate
- Improving Predictive Value of Poor Performance
- Root Cause Analysis of VTUAV Fire Scout's Nunn-McCurdy Breach
- Evaluating Solid Rocket Motor Industrial Base Consolidation Scenarios
- Managing Supply Chain Cyber Risks To DoD Systems and Networks
- Looking Back at PortOpt: An Acquisition Portfolio Optimization Tool
- Predicting the Effect of Schedule on Cost
- Recent Developments in the Joint Strike Fighter Durability Testing

Test and Evaluation: Statistical Methods for Better System Assessments
- Assessing Submarine Sonar Performance Using Statistically Designed Tests
- Applying Advanced Statistical Analysis to Helicopter Missile Targeting Systems
- Tackling Complex Problems: IDA's Analyses of the AN/TPQ-53 Counterfire Radar
- Improving Reliability Estimates with Bayesian Hierarchical Models
- Managing Risks: Statistically Principled Approaches to Combat Helmet Testing
- Validating the Probability of Raid Annihilation Test Bed Using a Statistical Approach

Technological Innovation for National Security
- Acquisition in a Global Technology Environment
- Lessons on Defense R&D Management
- Commercial Industry R&D Best Practices
- Strengthening Department of Defense Laboratories
- Policies of Federal Security Laboratories
- The Civilian Science and Engineering Workforce in Defense Laboratories
- Technology Transfer: DoD Practices

IDA | RESEARCH NOTES

4850 Mark Center Drive • Alexandria, VA 22311-1882
ida.org
@IDA_org