

Measuring Policy Effectiveness of Cyber Deterrence and Defensibility



22 March 2024

STRATEGIC MULTILAYER ASSESSMENT

Authored by: Mr. Jason Healey

Series Editor: Eric Kuznar, NSI Inc.

This paper was written for Strategic Multilayer Assessment's 21st Century Strategic Deterrence Frameworks project. For more information contact Mariah Yager at mariah.c.yager.ctr@mail.mil.

Mr. Jason Healey, Columbia University's School of International and Public Affairs

Jason Healey is a Senior Research Scholar at Columbia University's School for International and Public Affairs, specializing in cyber risk and conflict, and a part-time Senior Strategist at the National Risk Management Center at the US Cybersecurity and Infrastructure Security Agency. He has taught and mentored hundreds of students who have pursued careers at the White House, in the finance sector, civil society, and various other fields. Before joining Columbia University, he served as the founding director of the Cyber Statecraft Initiative at the Atlantic Council, where he established the global "Cyber 9/12" student cyber-policy competition. He edited the first history of conflict in cyberspace, titled "A Fierce Domain: Cyber Conflict, 1986 to 2012." Jason is a frequent keynote speaker on cyber risk and conflict, recognized as a "top-rated" speaker for the RSA Conference and recipient of the inaugural "Best of Briefing Award" at Black Hat.



Jason Healey played pivotal roles as a founding member of both the Office of the National Cyber Director at the White House (2022) and the world's first cyber command, the Joint Task Force for Computer Network Defense, established in 1998. In these positions, he contributed as an early pioneer of cyber threat intelligence. During a previous tenure in the White House, he served as a director for cyber policy, leading efforts to secure US cyberspace and critical infrastructure. He established Goldman Sachs' initial cyber incident response capability and subsequently oversaw the bank's crisis management and business continuity operations in Asia. He also served as the vice chair of the Financial Services Information Sharing and Analysis Center (FS-ISAC). Additionally, Jason Healey sits on the review boards of the DEF CON and Black Hat hacker conferences, contributed to the Defense Science Board task force on cyber deterrence, and previously served as president and founding board member of the Cyber Conflict Studies Association. He commenced his career as a US Air Force intelligence officer, holding positions at the Pentagon and the National Security Agency.

Measuring Policy Effectiveness of Cyber Deterrence and Defensibility

Mr. Jason Healey, Columbia University's School of International and Public Affairs¹

The United States needs more effective metrics to determine if integrated cyber deterrence is working (as called for in the 2023 cyber strategy for the Department of Defense (DOD)) as expected and to separate these effects from those of improving the overall defensibility of the Internet (as called for in the White House's National Cybersecurity Strategy). Fortunately, cyber deterrence is not like nuclear deterrence. Because of repeated interactions over time, it should be possible to measure to what degree deterrence is working to moderate the behavior of nation-state threat actors, not just—as with nuclear weapons—when it has failed.

This paper briefly examines the history of active cyber defense and deterrence, as well as earlier ideas on measuring effectiveness of cyber deterrence, before proposing new frameworks to measure if US government efforts at integrated deterrence and defensibility are succeeding at the strategic level. The main framework is based on relatively simple curves of adversary activity over time. In brief, measures to improve defensibility might be working if there is a downward trend (or decrease in slope) in the frequency and severity of general cybersecurity incidents. Such a decrease would likely tell us little about the success of integrated deterrence, which would require a downward trend (or decrease in slope) of frequency and severity of incidents by determined state threat actors.

If integrated deterrence is as successful a strategy as anticipated by DOD, the impact should be substantial enough to show up as a strong downward turn. Anything less may suggest that a strategy of integrated deterrence is insufficient and may need to be bolstered, supplemented with other strategies, or replaced. Without trendlines, DOD and policymakers cannot easily know. Though this approach may seem obvious, there have been few serious efforts, at least outside of classified channels.

¹ Contact Information: jh3639@sipa.columbia.edu

The new cyber strategy of the US Department of Defense (DOD) relies heavily on cyber capabilities as a part of integrated deterrence in the cyber domain (*Summary: 2023 Cyber Strategy of the Department of Defense*, 2023). Yet the United States must simultaneously deter four advanced adversaries: China, Russia, Iran, and North Korea. The DOD might need to either generate as much capability as all the nation’s adversaries combined or be four times more efficient or effective.

With so much reliance on integrated deterrence, the United States needs more effective metrics to determine if it is indeed working as expected and to separate these effects from those of improving the overall defensibility of the Internet, as called for in the new National Cybersecurity Strategy (*National Cybersecurity Strategy*, 2023, p. 1). Fortunately, cyber deterrence is not like nuclear deterrence. Because of repeated interactions over time, it should be possible to measure to what degree deterrence is working to moderate the behavior of nation-state threat actors, not just—as with nuclear weapons—when it has failed.

This paper briefly examines the history of active cyber defense and deterrence, as well as earlier ideas on measuring effectiveness of cyber deterrence, before proposing new frameworks to measure if US government efforts at integrated deterrence and defensibility are succeeding at the strategic level, convincing adversaries to limit their activity, or are only tactical wins, keeping adversary heads down only so long as “suppressing fire” is sustained.

The main framework is based on relatively simple curves of adversary activity over time. In brief, measures to improve defensibility might be working if there is a downward trend (or decrease in slope) in the frequency and severity of general cybersecurity incidents.² Such a decrease would likely tell us little about the success of integrated deterrence, which would require a downward trend (or decrease in slope) of frequency and severity of incidents by determined state threat actors.³

Curves like this serve two purposes. They provided a simple, idealized demonstration of the concepts involved, like a supply-demand curve or graph of GDP over time in an Economics 101 course. They can also illustrate the basic trends and impact of policy and operational decisions. Imagine the poor central banker trying to understand if an economy was entering a recession, or the impact of interest-rate changes in response, without a time-series graph of GDP over the last decade? This is, broadly, what cybersecurity professionals have been facing, often relying on “logic and analogy, rather than facts” (Maness et al., 2023).

² Measures to improve defensibility include ensuring software is more secure by design or that enterprises are implementing zero-trust architectures.

³ Measures to implement integrated deterrence include disruptive counter-cyber operations, sanctions, or arrests of nation-state cyber threat actors, or teaming with allies to publicly calling out adversary activity that violates global norms (or US national security).

To return to the economics example, decisionmakers may care less whether the exact rate of inflation rose from 6 percent to 7.8 percent or 8.5 percent; the crucial factor is that inflation is already high and rising. For most purposes, precision is less important for policymakers compared to understanding whether the magnitude and trends over time match expectations.

If integrated deterrence is as successful a strategy as anticipated by the DOD, the impact should be substantial enough to show up as a strong downward turn. Anything less may suggest that a strategy of integrated deterrence is insufficient and may need to be bolstered, supplemented with other strategies, or replaced. Without trendlines, the DOD and policymakers cannot easily know. Though this approach may seem obvious, there have been few serious efforts, at least outside of classified channels.

A Quick History of Active Cyber Defense and Deterrence

Since professional militaries are not culturally pre-disposed to passively wait for a blow which seems certain to fall, the DOD has accordingly explored concepts of deterrence and a more active defense against cyber-attacks. As early as 1996, a Defense Science Board (DSB) study recommended the department develop “rules of engagement for self-protection (including active response)” (*Report on the Defense Science Board Task Force on Information Warfare-Defense (IW-D)*, 1996, ES-10). Since then, new executive and legislative authorities have eased some of those constraints, so that military officials report that “as the adversary tries to maneuver, we can actually stay with the adversary” (Pomerleau, 2018).

Seen by some as perhaps even more important than active defense, cyber deterrence also has a long history. The same 1996 DSB study that called for a more active response also argued that “In the information age as in the nuclear age, deter is the first line of defense” (*Report on the Defense Science Board Task Force on Information Warfare-Defense (IW-D)*, 1996, ES-3). The Joint Staff went further just a few years later, wanting not just to deter adversaries from employing offensive capabilities against the United States, but implausibly from even establishing such relatively inexpensive, easily available, and inherently useful capabilities in the first place (*National Military Strategy for Cyberspace Operations*, 2006, p. 13). Deterrence, as it turns out, has its limits.

However, more recent DOD thinking seeks to capitalize on three related characteristics of cyberspace that might lead to more deterrent success. The first two were summarized in the statement in the new strategy that “cyber capabilities held in reserve or employed in isolation render little deterrent effect on their own” (*Summary: 2023 Cyber Strategy of the Department of Defense*, 2023). That is, unlike

nuclear weapons, their *deterrent value comes from their use*, and that they are unlikely to be able to deter on their own and *must be used with other instruments of national power*.⁴

The third characteristic is related to the “constant contact” mentioned earlier, or as the Department of Defense now expresses it, the need to “campaign in and through cyberspace” to “generate insights about cyber threats” and “advance Joint Force objectives” (*Summary: 2023 Cyber Strategy of the Department of Defense, 2023*). One major difference with nuclear and kinetic warfare, is that in cyber conflict the adversaries have countless, repeated interactions below the threshold of death and destruction, over many months, years, and even decades. These repeated interactions can both be used to reinforce deterrence (per the first point above) and track success, such as by the relative magnitude and frequency of adversary cyber operations over time. Cyber deterrence does not suffer the same fundamental measurability problems as nuclear deterrence, which could only be tracked by the lack of one major activity (launching a nuclear weapon at an adversary).

Measuring Deterrence in Cyberspace

A 2019 paper, co-written by the present author, proposed several such “rough-and-ready” metrics to suggest if deterrence was succeeding or not (Healey & Jenkins, 2019). The frameworks were necessarily rather simple, as so many cyber operations remain undetected, unreported, or have uncertain effects. Nearly five years later, there are few if any such metrics, at least reported in open source.

For some very small number of targets, the United States might demand absolutely zero incidents: There should never be any adversaries detected in the control rooms of nuclear power plants, for instance, or in nuclear command-and-control systems. If any such presence is detected, then deterrence might be said to have failed.

Most other metrics rely on trends over time. When announcing the White House’s new offensive policies in 2018, then-National Security Advisor John Bolton specified that incidents like China’s 2015 intrusion of the Office of Personnel Management were exactly “the kind of threat . . . from hostile foreign actors we’re determined to deter” (Bolton, 2018, September 20). In the simplest terms, if such OPM-magnitude incidents subsequently decrease, it is possible deterrence has worked; if they spike substantially higher, deterrence has likely not worked as expected.

⁴ Note: It has not historically always been the case that cyber deters only when combined with other instruments of national power. When discussing how to respond to Russia’s interference in the 2016 presidential elections, several White House participants reported being deterred by the threat of Russian escalation in cyberspace, either by attacking election systems directly (Ben Rhodes) or disruption of the US electrical grid (James Clapper). See *Not the cyber deterrence the United States wants*. Healey, J. (2018, June 11). <https://www.cfr.org/blog/not-cyber-deterrence-united-states-wants>.

So, while details of each incident are important, the major policy question around the success of deterrence is *whether threat actors are attacking more or less frequently over time and with higher or lower consequence*. Such trend lines cannot capture causation (for this, other sources, such as intelligence, are usually needed) but can be suggestive.

For example, after the agreement between presidents Obama and Xi in 2016 to limit espionage for commercial gain, FireEye (now Google Mandiant) reported a massive reduction in such operations (see Figure 1), a reduction later confirmed by executives from the National Security Agency and the Department of Justice (*Redline drawn: China recalculated its use of cyber espionage, 2016; Graff, 2018*). Chinese operations later bounced back to normal levels, but this only serves to highlight three issues:

1. Trend reporting provides useful insights.
2. These insights can suggest whether national-security policies seem to be working.
3. Adversaries will resume their activity unless repeatedly engaged by diplomacy, law enforcement, or military and intelligence power.

US Cyber Command’s goal to “to improve the security and stability of cyberspace” requires more than just the never-ending suppressing fire of persistent engagement (*Achieve and maintain cyberspace superiority: Command vision for US Cyber Command, 2018*).

It has been five years since the DOD announced their intention to defend forward and operate with agility and persistence, yet it does not seem that there has been any such reduction in adversary activity when compared to what was accomplished with the help of diplomacy against Chinese commercial espionage in 2016. Going from tactical engagement to tactical engagement—even when successful—was not enough to win in Vietnam, nor in Afghanistan. It is not clear why it might be more successful in cyberspace. The United States must deter at scale and over time.

Idealized Analysis on Deterrence and Defensibility

The rough-and-ready frameworks summarized above were useful when first suggested, several years ago, but must be extended further, to better differentiate between alternative explanations.

ACTIVE NETWORK COMPROMISES CONDUCTED BY 72 SUSPECTED CHINA-BASED GROUPS BY MONTH




Figure 1: Graph from cyber intelligence company FireEye showing reduction in China-based intrusion cases, from February 2013 to May 2016

The underlying dynamics can best be covered starting with just a few factors:

1. *Severity* is not just an objective measure of harm but is also influenced by subjective factors, especially the perception of US decisionmakers or the public. Several incidents (such as Chinese theft of information from OPM) caused more alarm than would have been suggested by their objective factors alone. Likewise, the severity of objectively similar campaigns—such as Russian ransomware after the invasion of Ukraine—will be higher during a geopolitical crisis. A more detailed analysis might split these factors out and assess them separately, but they remain grouped here for simplicity.
2. *Frequency* can be a simple count of incidents and campaigns in a certain period or a more qualitative measure, such as “increasing/decreasing” or a specified spectrum like “rare-uncommon-frequent-extremely frequent,” each with specific definitions.
3. *Intensity* is a function of the severity and frequency (the specifics of which—for example, is it an additive or a multiplicative function—can be determined later).⁵
4. *Time* can demonstrate how the trends change.

These factors can shed light on both defensibility and deterrence.

Defensibility (as called for in the US National Cybersecurity Strategy) represents a state where cyberspace is more defense friendly: It is difficult for typical threat actors to succeed in their goals and relatively easy for defenders to succeed at theirs (*Building a defensible cyberspace: Report of the New York Cyber Task Force, 2017*). Cyber incidents still occur, but they are unlikely to cascade or cause systemic problems. One of the innovations that most improved defensibility was Microsoft implementing automated patch updates in the 1990s, which substantially and nearly immediately simplified the process of securing computers (*Building a defensible cyberspace: Report of the New York Cyber Task Force, 2017*). More recently, programmers—encouraged by the White House—are increasingly moving to coding software in Rust and similar languages to almost entirely eliminate memory vulnerabilities, one of the most common and pernicious avenues of attack (*FACT SHEET: Biden-Harris Administration releases end of year report on Open-Source Software Security Initiative, 2024*).

For this paper, a defensibility metric would measure the implied success of efforts against cyber-attacks from all sources, though primarily cyber-crime (since this is the vast majority of all incidents).⁶

⁵ Note that to compute a function of frequency and intensity, both need to be the same kind of measurement. Most likely both would need to be ordinal (categorized and ranked). For reasons like these, the actual function of how to combine severity and frequency to develop a metric for intensity is left for further, more in-depth research.


⁶ There are only 429 incidents from 2000 to 2020 included in the main academic database of nation-state cyber incidents compared to nearly 5200 confirmed breaches in a single year, included in the main cybersecurity database. See Maness et

Deterrence (as called for in the DOD Cyber Strategy), by comparison, measures the US impact only against the very small subset of those incidents from the adversaries the United States wishes to deter, primarily Russia, China, North Korea, and Iran. It includes a wide set of measures, from diplomacy to on-net operations to disrupt and frustrate adversaries, to make them doubt the “belief that they can conduct unattributed coercive actions against the United States” (*Summary: 2023 Cyber Strategy of the Department of Defense, 2023*).

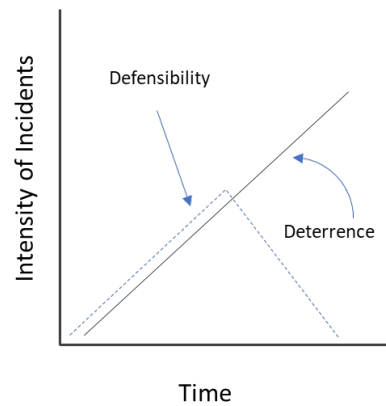
The US National Defense Strategy explained that integrated deterrence “entails developing and combining our strengths to maximum effect, by working seamlessly across warfighting domains, theaters, the spectrum of conflict, other instruments of U.S. national power, and our unmatched network of Alliances and partnerships” (*2022 National Defense Strategy, 2022, p. 1*). The subsequent cyber strategy did not particularly expand on how integrated deterrence applied to cyberspace.

Though related, defensibility and deterrence are not likely to be tightly correlated. Because nation-state adversaries are persistent and capable, improvements in defensibility are unlikely to have much impact on their operations. Likewise, targeted deterrence against high-end cyber adversaries is unlikely to affect most cyber criminals.

Imagine stylized curves, such as those for GDP growth or interest rates over time in an introductory economics course, represented below as Cases 1 through 4. These compare intensity of cyber campaigns over time.



Case 1: Nether Defensibility or Deterrence Appear to be Working



Case 2: Defensibility Working, No Change to Deterrence


al. (2023) (for the academic dataset) and the *Verizon Data Breach Investigations Report (2023)* (registration is required for the latter resource).

Case 1 illustrates the base case with the curves for both defensibility and deterrence heading up and to the right if the actions called for in the National Cybersecurity Strategy and DOD Cyber Strategy are no more effective than all other past strategies.

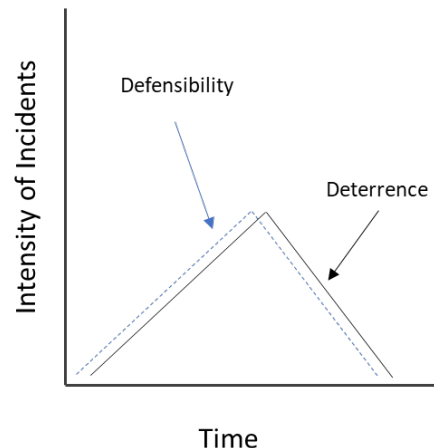
By contrast, Case 2 demonstrates the shape of the curves if there is an increase in defensibility, leading to a decrease in general cybersecurity incidents. However, nation-state adversaries remain undeterred and able to continue their depredations.

The reverse is true in Case 3, where efforts by the US government to deter adversaries have succeeded but with no corresponding decrease in overall incidents, as with the top-right quadrant. The decrease in deterrence may of course be correlated with other factors, such as a decline in global crises and conflict, so that adversaries unilaterally reduce their activities. Fortunately, these factors should be visible and measurable, allowing analysts to assess how decline in adversary activity might be due to integrated deterrence.

A more difficult factor in determining if deterrence led to any declines in adversary activity over time is visibility. Any decrease may only be because adversaries have gotten harder to detect. While this is a core issue with any strategy seeking to deter cyber adversaries, such observability is less of an issue for this approach. Assuming policymakers' priorities are indeed to reduce the most brazen, reckless, and damaging cyber activity, these campaigns are usually the most obvious over time. Moreover, it may be a safe assumption that over the medium term, visibility is relatively constant; that is, while defenders' momentary visibility into adversary activity may wax or wane, it may stay close to an average over time.



Case 3: Deterrence Working but Not Defensibility



Case 4: Both Defensibility and Deterrence Working

Case 4 is the happy situation of major decreases in both curves suggesting that the actions of both the White House and DOD strategies have been successful. Again, there will be a range of other plausible causes, but many of these can be identified and controlled for, but only if the trend is first established.

An even simpler representation, which still can provide policy insights, uses just a 2x2 matrix, as shown in Table 1.

Table 1. Incidents Frequency and Severity

	Frequency of Incidents Increases	Frequency of Incidents Decreases
Average Severity Increases	Worst quadrant: Nothing the United States does is working. (Similar to Case 1).	Second worst quadrant: Even though overall incidents decrease, most impactful incidents continue to worsen. (Similar to Case 2).
Average Severity Decreases	Second-best quadrant: Even though average severity is decreasing, incidents in general are more frequent. (Similar to Case 3).	Best quadrant: Everyone takes credit. (Similar to Case 4).

From Idealized Analysis to Policy and Operational Relevance

But can it work in practice?

Several commercial and academic datasets could be used to test these hypotheses, as could classified datasets held by the DOD and the intelligence community. Any kind of exact measurement of cyber incidents is impossible, as too many incidents go undetected; however, the particulars matter far less than the magnitude and direction of the curve.

Plotting the curve for defensibility might be straightforward using datasets in the private and public sectors. In the private sector, the most obvious source is the VERIS Community Database, the basis for the gold standard of reports of cyber intrusions, the Verizon Data Breach Investigations Report (VDBIR) (*Verizon launches the VERIS Community Database, 2023*). One of the best commercial data providers is Advisen, whose dataset includes more than 90,000 incidents and is heavily used by the insurance industry, financial institutions, and researchers (*Cyber loss data, n.d.*). Major technology companies, cybersecurity providers, and intelligence companies—such as Microsoft, Google, Recorded Future, and

CrowdStrike—might be persuaded to assist through partnerships such as the Joint Cyber Defense Collaborative with the Cybersecurity and Infrastructure Security Agency (CISA).

Within government, CISA is building out datasets on cyber incidents and risk, efforts which will become substantially more comprehensive once they start receiving incident reports mandated by the recent CIRCIA act (*Cyber Incident Reporting for Critical Infrastructure Act of 2022 (CIRCIA)*, n.d.). Unfortunately, CISA may not be regularly assigning severity (such as by using the National Cyber Incident Scoring System) to all incidents in their data set (Member of Cybersecurity and Infrastructure Security Agency, personal communication, October 17, 2023; *CISA National Cyber Incident Scoring System (NCISS)*, 2020).

The deterrence curve will be harder to plot, as information on adversaries is held more closely. Moreover, because each adversary is different, each would likely need their own chart plotting intensity over time. However, such information does exist—sometimes in time-series format—within cyber-intelligence and response companies, major technology vendors, and of course within parts of the US intelligence community and military. The academic Dyadic Cyber Incident Database, which tracks 429 nation-state incidents from 2000 to 2020, shows an annual decrease in the number of incidents since roughly 2017 (Maness et al., 2023). However, this decrease does not necessarily mean that deterrence is working, as this measurement does not include severity.

Beyond sufficient data, several other issues would need to be addressed. First, these charts only suggest correlation and certainly not causation. But this is the case with any time-series, so responsible analysts must explore alternate hypotheses using other data and case studies.

Fortunately, the proponents for integrated deterrence suggest that its impact should be especially marked: As a previous national security advisor put it, the president “authorized offensive cyber operations [...] not because we want more offensive operations in cyberspace, but precisely to create the structures of deterrence that will demonstrate to adversaries that the cost of their engaging in operations against us is higher than they want to bear” (Bolton, 2018). Accordingly, anything less than a strong decrease (that is, a trend with substantial negative slope, as in Cases 3 and 4), suggests deterrence is not working as expected.

Second, statisticians and data scientists will need to review the data to ensure that it is properly processed. For example, economic analysts do not just observe whether a company’s stock has risen or dropped, but assess the cumulative abnormal returns, to determine a better appropriate baseline (see, for example, Knight & Pretty, 2001). Similar normalization will be needed to determine cyber baselines and trendlines.

Third, research and investment are needed to define severity and assign scores to thousands of incidents and campaigns. Severity is difficult to assess, as there are problems with any chosen method. Insurance companies might be fine with a financial impact measured in dollars, but that is less useful

for determining impact on critical infrastructure, degradation to national security, or national outrage demanding a military response.

Any fine-grain qualitative categorization (such as the 10-point scale of the Dyadic Cyber Incident Database) can be overly arbitrary and can have counting problems. For example, Russia's massive SolarWinds intrusion campaign, which actively exploited over 100 targets including government agencies (and infected up to 18,000 others) would probably be counted only once and might rate only a 4 out of 10 in severity. Despite causing one of the most massive incident response efforts ever, and directly leading to US sanctions of Russia (*How U.S. cyber policy changed after SolarWinds*, 2021), the incident would likely be coded no higher than any other case of cyberespionage (Maness et al., 2022). Perhaps the easiest successful solution is just three logarithmic distinctions of severity: minimal and local, moderate and national-security relevant, and truly massive.

Fourth, investment and research are needed to better understand correlation and causation. For example, US policymaking is more likely to be the cause for improved defensibility if there appears to be a continued rise of incidents in other advanced nations, such as Japan, which did not take such measures.

Conclusion

Modern economics is unthinkable without time-series charts to guide decisions, like whether to raise interest rates based on whether GDP or inflation are shrinking or growing. But cybersecurity is nowhere near as observable or analytically mature as economics. The above idealized charts can be useful—even without including any actual data—to understand the basic principles of the policy expectations.

To extend the economics analogy one step further, cybersecurity—including in the DOD—is largely at the micro level, examining the security of individual systems or organizations, and not the equivalent of macroeconomics, focused on the dynamics of the systems as a whole.

Charts tracking these macro issues will not be easy to develop but must be accomplished to determine if cyberspace is becoming more defensible and if adversaries are deterred from conducting cyber campaigns as a result of our integrated deterrence strategies. The needed investment to determine if the DOD is meeting the expectations of itself for integrated deterrence may cost several million dollars—but would be substantially cheaper than building and employing cyber capabilities. It is an investment worth making.

References

- 2022 National Defense Strategy of the United States of America*. (2022). U.S. Department of Defense. <https://media.defense.gov/2022/Oct/27/2003103845/-1/-1/1/2022-NATIONAL-DEFENSE-STRATEGY-NPR-MDR.PDF>
- Achieve and maintain cyberspace superiority: Command vision for US Cyber Command*. (2018). US Cyber Command. <https://www.cybercom.mil/Portals/56/Documents/USCYBERCOM%20Vision%20April%202018.pdf?ver=2018-06-14-152556-010>
- Bolton, J. (2018, September 20). *Transcript: White house press briefing on national cyber strategy - Sept. 20, 2018* [Briefing Transcript]. *The White House*. <https://news.grabien.com/making-transcript-white-house-press-briefing-national-cyber-strateg> Video available here (around minute 22:30): <https://www.c-span.org/video/?451807-1/national-security-adviser-john-bolton-cyber-strategy-audio-only>
- Building a defensible cyberspace: Report of the New York Cyber Task Force*. (2017). Columbia University, School of International and Public Affairs. https://www.sipa.columbia.edu/sites/default/files/2022-09/3668_SIPA%20Defensible%20Cyberspace-WEB.PDF
- CISA National Cyber Incident Scoring System (NCISS)*. (2020, September 30). Cybersecurity and Infrastructure Security Agency. <https://www.cisa.gov/news-events/news/cisa-national-cyber-incident-scoring-system-nciss>
- Cyber loss data*. (n.d.). Advisen. <https://www.advisenltd.com/data/cyber-loss-data/>
- Cyber Incident Reporting for Critical Infrastructure Act of 2022 (CIRCI A)*. (n.d.). Cybersecurity and Infrastructure Security Agency. <https://www.cisa.gov/topics/cyber-threats-and-advisories/information-sharing/cyber-incident-reporting-critical-infrastructure-act-2022-circia>
- FACT SHEET: Biden-Harris Administration releases end of year report on Open-Source Software Security Initiative*. (2024, January 30). The White House. <https://www.whitehouse.gov/oncd/briefing-room/2024/01/30/fact-sheet-biden-harris-administration-releases-end-of-year-report-on-open-source-software-security-initiative/>
- Graff, G.M. (2018, October 11). How the US forced China to quit stealing—Using a Chinese spy. *Wired*. <https://www.wired.com/story/us-china-cybertheft-su-bin/>
- Healey, J. (2018, June 11). Not the cyber deterrence the United States wants. *Council on Foreign Relations*. <https://www.cfr.org/blog/not-cyber-deterrence-united-states-wants>
- Healey, J., & Jenkins, N. (2019, 28-31 May). *Rough-and-ready: A policy framework to determine if cyber deterrence is working or failing* [Conference paper]. 11th International Conference on Cyber Conflict, Tallinn, Estonia. https://ccdcoe.org/uploads/2019/06/Art_07_Rough-and-Ready.pdf
- How U.S. cyber policy changed after SolarWinds*. (2021, July 4). CBS News. <https://www.cbsnews.com/news/solarwinds-60-minutes-2021-07-04/>
- Knight, R., & Pretty, D. (2001). *Reputation & value: The case of corporate catastrophes*. Oxford Metrica. <http://www.oxfordmetrica.com/public/CMS/Files/488/01RepComAIG.pdf>
- Pomerleau, M. (2018, November 27). *Defense officials taking advantage of new cyber authorities*. Fifth Domain. <https://www.c4isrnet.com/dod/cybercom/2018/11/27/defense-officials-taking-advantage-of-new-cyber-authorities/>

-
- Redline drawn: China recalculated its use of cyber espionage.* (2016). Mandiant.
<https://www.fireeye.com/content/dam/fireeye-www/current-threats/pdfs/rpt-china-espionage.pdf>
- Reuters. (2018, November 8). *US accuses China of violating bilateral anti-hacking agreement.* CNBC.
<https://www.cnbc.com/2018/11/09/us-accuses-china-of-violating-bilateral-anti-hacking-agreement.html>
- Manness, R.C., Valeriano, B., Hedgecock, K., Jensen, B.M., & Macias, J.M. (2022). *Codebook for the Dyadic Cyber Incident and Campaign Dataset (DCID) Version 2.0* [Dataset]. drryanmaness.wixsite.com
<https://drryanmaness.wixsite.com/cyberconflict/cyber-conflict-dataset>
- Manness, R.C., Valeriano, B., Hedgecock, K., Jensen, B.M., & Macias, J.M. (2023). Expanding the Dyadic Cyber Incident and Campaign Dataset (DCID): Cyber Conflict from 2000 to 2020. *Cyber Defense Review*, 65-89.
https://cyberdefensereview.army.mil/Portals/6/Documents/2023_Summer/Maness_Macias%20et%20al%20CDR%20V8N2%20Summer%202023.pdf?ver=iJDHRpDvaq-hW4Zde6EwUg%3d%3d
- National Cybersecurity Strategy.* (2023). The White House. <https://www.whitehouse.gov/wp-content/uploads/2023/03/National-Cybersecurity-Strategy-2023.pdf>
- National Military Strategy for Cyberspace (U).* (2006). United States Office of the Chairman of the Joint Chiefs of Staff. <https://www.hsdl.org/?abstract&did=35693>
- Report of the Defense Science Board Task Force on Information Warfare-Defense (IW-D).* (1996). Office of the Undersecretary of Defense for Acquisition & Technology. <https://apps.dtic.mil/sti/pdfs/ADA319571.pdf>
- Summary: 2023 Cyber Strategy of the Department of Defense.* (2023, September 12). U.S. Department of Defense. https://media.defense.gov/2023/Sep/12/2003299076/-1/-1/1/2023_DOD_Cyber_Strategy_Summary.PDF
- Verizon data breach investigations report.* (2023). Verizon.
<https://www.verizon.com/business/resources/reports/dbir/2023/summary-of-findings/>
- Verizon launches the VERIS Community Database.* (2023, July 25). Verizon.
<https://www.verizon.com/about/news/verizon-launches-veris-community-database>

