THE CYBER DEFENSE REVIEW
2026, VOL. 11, NO. 2, 235-245
https://doi.org/10.55682/cdr/kjm8-q2dm

SENIOR LEADER PERSPECTIVE

We Are the Attack Surface: Conway's Law, the Sociotechnical Layer, and the Resilience the Next Decade Demands

Cmdr. Ryan P. Hilger, Ph.D.

U.S. Navy, USA

The defense cyber community has made genuine progress at the technical layer over the past decade. Continuous software delivery and security practices, supply chain security certification frameworks, and the maturation of dedicated military cyber forces represent real and consequential structural advances. That progress now reveals the layer above it. The seams adversaries exploit with consistent effectiveness are not technical seams—they are organizational seams, the predictable product of how complex work is organized across institutional boundaries at scale. This essay synthesizes a decade of research and acquisition practice to advance three connected arguments. First, existing cybersecurity and resilience frameworks systematically address the technical layer while leaving the sociotechnical and ecosystem layers under-addressed—a gap that represents the defining strategic liability of the coming decade. Second, Conway's Law and panarchy theory together explain why this gap persists: organizations produce systems that mirror their communication structures, and cross-scale dynamics ensure that fast-cycle compromises can cascade upward to destabilize strategic command and control. Third, addressing this gap requires deliberate attention across the full lifecycle of cyberphysical systems—from development through deployment, fielding, active defense, and sustainment—not only at initial program authorization. Artificial intelligence amplifies both the consequences of the problem and the cost of its continued deferral, while simultaneously offering new analytical tools for ecosystem-level situational awareness. The path forward requires treating resilience as a continuous lifecycle obligation, mandating ecosystem-level threat modeling, and recognizing operational commanders as the essential demand signal for resilient outcome specifications.

Keywords: Conway's Law; panarchy; cybersecurity; cyber resilience; defense acquisition ecosystems, cyberphysical systems

Commander Ryan Hilger is the Principal Assistant Program Manager for Submarine Launched Unmanned Undersea Systems in the Navy's Advanced Undersea Systems Program Office. He holds a doctoral degree from Colorado State University, where his dissertation examined cyber resilience frameworks across the organizational and ecosystem levels of defense acquisition systems. His published research includes work on cyber resilience in complex sociotechnical ecosystems, applications of machine learning to elicit shared contextual information, and use of Bayesian methods to predict non-transitive trust in cyber ecosystems.

THE SEAM IS THE VULNERABILITY

A vulnerability disclosure arrives on a Tuesday morning. The affected component is widely used across the defense industrial base—software managing workflow and documentation for dozens of programs. Within hours, exploitation attempts begin. By the end of the week, they number in the tens of thousands.

At a program with a mature software factory, the response is almost elegant. The team identifies the affected libraries, issues a patch through their continuous integration pipeline, and pushes an update within seventy-two hours. The investment in continuous delivery paid off exactly as intended. The continuous authorization-to-operate process catches the change cleanly. The cyber force is postured and watching. At the technical layer, everything works.

And yet, the same software platform runs on contractor networks two and three tiers down the supply chain—networks that share development pipelines with the prime. The patch requirement travels through a contract actions that may take months to formally deliver, let alone complete the patch, certify the build, and re-deliver to the government. The operational technology embedded in a deployed platform—a system designed for a long service life and updated on a scheduled depot cycle—may have no patch path at all. The seam was never in the code. It was in the organization.

This is not a hypothetical. It is a structural description of how the defense acquisition ecosystem operates, derived from a decade of research, program experience, and empirical analysis of how complex organizations actually build and sustain cyberphysical systems. The seams that adversaries exploit with consistent effectiveness are not accidents of negligence. They are the natural product of how humans organize complex work at scale—a dynamic that software engineer Melvin Conway (1968) first described and that defense acquisition has validated, with compound interest, in the decades since. My perspective on this problem is shaped by nearly a decade of working across major defense acquisition programs, where the organizational geometry these pages describe is an observable feature of every program ecosystem I have supported. The argument that follows reflects both that experience and a sustained program of research into the dynamics that make complex defense systems fragile.

The defense cyber community has spent the past decade doing what it set out to do: advancing the technical fundamentals. Continuous software delivery and security practices have materially changed how major programs develop and deliver code. The Cybersecurity Maturity Model Certification (CMMC) framework has begun raising the floor of the defense industrial base. United States Cyber Command and the service cyber components have matured into serious operational forces with genuine deterrent and defensive capabilities. Continuous authorization-to-operate processes have replaced static point-in-time reviews across an increasing number of programs. This is genuine progress. It is worth naming clearly, because it changes what comes next.

Technical-layer progress does not merely leave the problem partially solved. It reveals the layers above it. As endpoints and networks became harder targets, adversaries adapted. The seams that remain—and that are now exploited with systematic sophistication—are organizational seams, process seams, cross-scale seams. They are the natural geometry of how humans build complex things together at the pace and scale that defense programs demand. The compliance frameworks, certification requirements, and authorization processes that produced real technical-layer gains were not designed to address this layer. None of them were.

The stance this essay takes is corrective, not merely diagnostic: the technical-layer progress of the past decade was necessary and right, and it now positions the defense cyber enterprise to address the harder sociotechnical and ecosystem problem with the same deliberateness it brought to the technical layer over the last decade.

This essay advances three connected arguments. First, our frameworks have systematically addressed resilience (though I take issue with the usage of that term) at the technical layer while leaving the sociotechnical and ecosystem layers under-addressed. That gap is now the defining strategic liability. Second, Conway's Law and panarchy theory together provide the conceptual tools to close it, and the defense research base has already begun applying them. Third, this work must happen across the full lifecycle of cyberphysical systems—from development through deployment, fielding, defense, and sustainment—not only at authorization. Artificial intelligence raises the stakes of getting this right and the cost of continuing to defer it. The path forward is not mysterious. We know enough to move.

A DECADE OF TECHNICAL PROGRESS—AND THE LAYER IT REVEALS

Let me be precise about what the decade produced, because precision matters here. The software factory model demonstrated that significant portions of the defense industrial base can achieve continuous integration, delivery, and authorization at speeds that were impossible a decade ago. That is a real and consequential capability shift. The CMMC framework, however contested its rollout, established that cybersecurity practice in the defense industrial base is a contractual obligation, not a voluntary aspiration. The growth of operational cyber forces has given combatant commanders tools that simply did not exist in 2015. These are structural changes, not incremental improvements.

And yet a Defense Science Board task force concluded in 1987 that the major problems facing military software development were not technical problems but management problems (USD(A&S) 1987). A second task force reached the same conclusion in 2018, in nearly identical language, about programs that had been continuously reformed in the intervening thirty years (USD(R&E) 2018).

The acquisition community had not been standing still; it had produced frameworks, directives, and certification programs by the dozens. But the fundamental conditions had not changed.

Recent research examining the landscape of cybersecurity and resilience guidance documents reveals a complementary finding. An algorithmic analysis of 37 cybersecurity frameworks and 12 resilience frameworks drawn from outside the cyber domain—ecology, organizational theory, infrastructure resilience, public health—finds that no single existing guidance document comprehensively addresses cyber resilience across the levels of an organization or ecosystem (Hilger 2025). The frameworks governing defense programs (i.e., NIST) cluster at the technical level: protecting endpoints, securing networks, and managing vulnerabilities. The organizational and sociotechnical levels—where humans, processes, incentive structures, and the cultural histories of the acquisition enterprise actually shape system behavior—are systematically under-addressed. This is not a coincidence. The guidance we have reflects the organizational fragmentation it was written to remedy.

The distinction that matters here is between security, robustness, and resilience (Kott and Linkov 2019). Cybersecurity seeks to prevent adverse events and, when they occur, restore the system to its prior state. Robustness controls the depth of the loss when an event occurs—how much capability degrades under attack. Resilience encompasses both and adds a third and defining property: adaptation. A resilient system does not merely recover to the status quo ante. It reorganizes, restructures, or finds alternate means of functioning, emerging from adversity with at least equivalent capability and often greater. It is an emergent behavior in a complex adaptive system, which means people are involved, not just the technical system. The technical-layer investments of the past decade have substantially advanced security. Robustness has improved in programs that adopted continuous delivery. Adaptation—the property that distinguishes resilience from robustness, and the property that determines whether a system remains operationally viable under sustained adversary pressure—requires operating at the sociotechnical and ecosystem levels that our frameworks have not reached.

Critically, adaptation cannot be decomposed into a checklist item, specified as a discrete feature in a requirements document, or verified in an authorization-to-operate review. It emerges from the way people, processes, and technology interact under stress. This is precisely why compliance-oriented frameworks have not produced it. But emergence does not mean unspecifiable. The outcomes resilience must produce can be specified: What minimum critical functionality must this system maintain during a cyberattack? What degraded-mode behaviors must operators execute? Within what recovery timeline must the system restore full capability for the operational commander's mission to remain viable? These are acquirable requirements. They are also requirements that the acquisition system has not been systematically asked to produce. That is the work of the next decade.

CONWAY'S LAW, PANARCHY, AND THE ECOSYSTEM ADVERSARIES READ

Conway (1968) observed that organizations that design systems are constrained to produce designs that mirror their own communication structures (i.e., organizational chart, contracting relationships, etc.) and noted explicitly that the principle applied to military systems as well as commercial software. Defense acquisition professionals who have spent time in major weapon system programs will recognize the phenomenon without needing the academic framing.

The most vivid recent illustration is the relationship between the F-22 Raptor and the F-35 Lightning II. Both aircraft are Lockheed Martin products. Both are governed by the same interoperability requirements. Neither can communicate directly with the other—their data links are incompatible. In 2021, the Air Force executed Project Hydra to bridge the gap, connecting the two aircraft through a third platform, the U-2 Dragon Lady, acting as an airborne relay (Trevithick and Rogoway 2021). The organizational seam between the F-22 and F-35 programs became a system seam, which became a capability seam, which, in a contested environment, is a vulnerability seam. This is not an anomaly. It is the norm. It is Conway's Law expressed in kinetic terms.

The Submarine Warfare Federated Tactical System, a commercial off-the-shelf integration effort begun in 1997, today represents the synchronized output of more than a dozen program offices (Mitchell 2014). Its architecture is, in meaningful ways, a map of its organizational history. The seams between program offices are legible in its system interfaces. This is not a failure of management. It is the predictable geometry of complex work organized across institutional boundaries at the defense acquisition scale.

Conway's Law describes the mechanism. Panarchy theory describes the cross-scale dynamics that determine its strategic consequences (Allen et al. 2014; Holling 2001). Panarchy holds that complex adaptive systems—whether ecological, organizational, or sociotechnical—operate simultaneously across multiple nested scales, each with its own cycle speed and adaptive capacity. Fast, small cycles are nested within slower, medium cycles, which are in turn nested within the slowest, large cycles. In the defense acquisition ecosystem, technical patches and software updates represent the fast small cycle. Program structures, organizational boundaries, and contractor relationships represent the medium cycle. Doctrine, acquisition policy, and industrial base architecture represent the slow large cycle.

Two mechanisms in panarchy theory are directly relevant to cyber resilience. The first, which Holling (2001) called "remember," describes how the larger, slower cycle enables and constrains the faster smaller cycles within it. Continuous delivery capability at the technical layer—however mature—operates within an organizational medium cycle it cannot itself reform. The software factory can patch faster than the contracting vehicle can communicate the requirement downstream. The second mechanism, "revolt," describes how a fast small-scale disruption can cascade upward to destabilize the larger slow cycle. A supply chain compromise in a third-tier contractor network—a fast, localized event—can propagate through program ecosystem dependencies to affect the data fabric of a joint Command and Control architecture. SolarWinds was not a failure of technical-layer defenses. It was a panarchic revolt event: a fast compromise that cascaded upward through the medium cycle of organizational relationships to threaten the slow cycle of strategic Command and Control. The frameworks we had were designed for the fast cycle. Cross-scale dynamics were outside their aperture.

The strategic implication warrants stating plainly. Adversaries do not need zero-day exploits to map our attack surface. They can read organizational charts, contract award announcements, program office structures, and acquisition documentation. China's systematic approach to cyber operations—patient, methodical, ecosystem-aware—has been reading that documentation for decades, conducting the reconnaissance and resource development that sophisticated operations require long

before a weapon system reaches operational deployment (O'Neill 2022). The map is not classified. We publish it.

Expanding the aperture from the individual system to the program ecosystem makes the scale of this legibility concrete. A defense acquisition program is embedded in a surrounding ecosystem of government workforce, contractor organizations, second- and third-tier suppliers, enterprise infrastructure on which development occurs, open-source repositories from which code may be drawn, and the communities in which the human workforce lives and works. Each node in that ecosystem and each flow between nodes is a potential vector for adversary reconnaissance and exploitation. Ecosystem-level threat modeling reveals this. System-centric threat modeling, focused on the technical platform and its immediate users, does not.

Cherns (1976) observed that most engineers and managers acknowledge they are designing a sociotechnical system, yet bring substantial knowledge to the technical side and little to the social side. This is not human factors engineering or human systems integration. The consequences of this bias manifest in ways that will do as much to thwart as to facilitate the system's intended functions, as interdependencies grow in evolutionary, often invisible ways. The defense acquisition community is, decades later, living that observation—now amplified across an interconnected ecosystem of cyberphysical systems whose broader social architectures were never designed with adversarial exploitation in mind.

The reframe that follows is the central argument of this essay. If the attack surface is inherent—if Conway's Law ensures that organizational seams become system seams, and panarchic cross-scale dynamics ensure that fast-cycle compromises can cascade upward to destabilize strategic command and control—then the acquisition community's job is not to eliminate the attack surface. It is to understand it at the ecosystem level, specify what must survive across it under adversary exploitation, and build organizations capable of adapting when it fails. That is a different job description than the one compliance frameworks currently assign.

THE LIFECYCLE GAP—FROM DEVELOPMENT TO SUSTAINMENT

In June 2022, a critical vulnerability in Atlassian's Confluence platform was publicly disclosed (Atlassian 2022). The platform is widely used across the defense industrial base to manage software development. Within hours, adversary exploitation attempts began. By the end of the first day, they reached 20,000. Within weeks, they peaked at 100,000 attempts per day. A defense program with mature continuous delivery capability responded in days. A program managing operational technology embedded in a deployed platform may have been staring at a seemingly unsolvable patching problem. The technical fast cycle and the fielded cyberphysical system exist in different panarchic worlds.

That gap is not a technical failure. It is a lifecycle failure, and it is structural. The seam between the technical fast cycle and the fielded platform is one of the widest and most persistently under-addressed seams in the defense acquisition ecosystem. The organizational structure of the sustainment enterprise—separate from the development enterprise, governed by different contracts, staffed by different workforces, operating on different timelines—determines the response tempo. That tempo is a security parameter. Most programs have never explicitly specified or measured it, though a few might be measuring Mean Time to Patch.

Walking the full lifecycle spine makes the compounding nature of this problem visible. At the development phase, the primary seams are between the government program office and the prime contractor, and between the prime and its second- and third-tier suppliers—Conway's Law ensures the architecture reflects those boundaries. At deployment, the seam is between the development enterprise and the operational force: the documentation, training, and configuration management that carries the system from the contractor's facility to the operator's hands. At fielding, legacy platforms are integrated into new architectures; each integration point is a seam between organizational histories. At the defense phase—the sustained operation of fielded systems—the seam is between the cyber force that can detect and respond and the program office and contractor that hold the authority and capability to remediate. At sustainment, the seam is between the organic capability of the program and the depot and contractor relationships that govern long-term maintenance. Each transition compounds the vulnerabilities of the last.

Cyberphysical systems—ships, aircraft, vehicles, installations, industrial control systems—are where these lifecycle seams are widest and the stakes are highest. These systems may have decade-plus operational lives. They are the systems where compromise has physical consequences. The failure of a network is a cybersecurity event. The failure of a cyberphysical system is a safety event, an operational readiness event, and a potential mission failure. As Secretary of the Navy Thomas Modly stated in his keynote address at WEST Conference (2019), "…our vulnerabilities may be so great that we may not even be able to get our ships underway out of San Diego if they hit us a certain way".

The authorization to operate asks whether a system is secure enough to begin operating. It does not ask what the system must do when compromised mid-operation: what degraded-mode behaviors the operational commander requires, what minimum critical functionality must be preserved during an active attack, or within what timeline the system must recover to restore mission capability. These are the resilient outcome specifications that the acquisition system has not been systematically asked to produce. They are the difference between an acquisition community that manages compliance and one that manages operational risk jointly with the warfighter across the full lifecycle.

The practical implication is direct: resilience cannot be retrofitted at authorization. The decisions that determine a system's resilience—its software architecture, its supplier dependencies, its update mechanisms, its degraded-mode design, the organizational boundaries that govern its integration and sustainment—are made years before the system reaches the operator. The operator inherits whatever was built in. Changing operational outcomes requires changing acquisition decisions, which in turn requires policy and doctrine that treat resilience as a lifecycle obligation rather than a checkpoint at the beginning of service.

ARTIFICIAL INTELLIGENCE—AMPLIFIER, ACCELERANT, AND OPPORTUNITY

Artificial intelligence (AI) does not change Conway's Law. It amplifies its consequences across all scales of the panarchic hierarchy—and it does so in both directions.

AI as Amplifier of Ecosystem Vulnerability

The Joint All-Domain Command and Control (JADC2) vision places artificial intelligence-enabled decision support at the center of joint command and control (DoD 2022). Intelligent algorithms will

ingest data from the federated data fabric, produce knowledge superiority outcomes, and support commanders at multiple echelons at machine speed. This is a compelling operational capability. It is also a significant concentration of risk at precisely the seam where the panarchic hierarchy is most fragile—the interface between the medium cycle of organizational structure and the large slow cycle of strategic command.

If the ecosystem feeding the data fabric is compromised—through supply chain infiltration, data poisoning, lateral movement from a contractor network, or any of the attack pathways the program ecosystem model reveals—AI algorithms will ingest corrupted inputs and produce corrupted outputs at machine speed, below human detection thresholds, and potentially in ways that subtly shift commander decision-making rather than triggering obvious system failures. It may be similar to the responses from Iranian scientists and engineers during Stuxnet, just faster and even more sophisticated now. An adversary who can influence what the algorithm sees can influence what the commander decides, without ever appearing to have acted. This is a panarchic revolt event operating at machine speed. It is not a future threat. It is a present design condition of any artificially intelligent system built on the acquisition ecosystem we have today.

AI as Accelerant of Adversary Exploitation

The tempo gap between vulnerability disclosure and acquisition system response was already strategically dangerous before AI entered the picture. Adversaries equipped with these tools can conduct the reconnaissance, resource development, and initial access activities of ecosystem exploitation faster, at a greater scale, and with greater precision than before. Mapping organizational structures, identifying key personnel and supplier relationships, synthesizing open-source intelligence about program architectures—these are the pattern-recognition and data-synthesis tasks at which machine learning excels. The ecosystem's legibility, which Conway's Law and panarchy produce at every scale simultaneously, becomes more exploitable when adversaries have tools to read it comprehensively.

China's approach to cyber operations has demonstrated decades of patience and systematic ecosystem analysis. These tools accelerate that analysis without changing its fundamental character. An exposure window that was operationally problematic five years ago may be unacceptable today—and the window will narrow further as adversary capabilities mature. The lifecycle gap is not a static problem. It is deteriorating.

AI as Opportunity to Enhance Ecosystem-Level Resilience

The same capabilities that make these tools dangerous in the hands of adversaries make them potentially transformative for defense. Machine learning can map program ecosystem dependencies at a scale and granularity that human analysis cannot match—identifying critical seams, modeling risk propagation pathways across panarchic scales, flagging anomalies in supply chain behavior, and giving acquisition leaders genuine visibility into their attack surface for the first time. Research applying algorithmic and natural language processing methods to cybersecurity framework analysis has demonstrated the potential of this approach at the policy level; the same logic applied to ecosystem threat modeling could transform program-level situational awareness (Hilger 2025).

AI also has a role in the resilient outcome specification problem. If adaptation is an emergent property that cannot be directly specified in a requirements document, simulation and wargaming

tools can model system behavior under adversarial stress, identify the organizational and technical conditions that produce resilient outcomes, and help acquisition teams design toward those conditions. This is an underexplored capability development opportunity that the defense community should pursue deliberately in the coming decade. Additionally, AI may be able to help us write better requirements that better enable resilient outcomes and to design the organizational and contracting structures to achieve them.

The net assessment is sobering but not defeatist. AI does not create the attack surface—Conway's Law and organizational complexity do that. It raises the stakes of getting the ecosystem resilience problem right and the cost of continuing to defer it. It also offers, for the first time, analytical tools adequate to the scale of the problem.

A WAY FORWARD—THREE LANES

The argument this essay makes is corrective and diagnostic. Diagnoses are only useful when they point toward action. What follows is a set of directional priorities organized by the three audiences the problem most directly implicates.

For Policymakers

Establish cyber resilience as a continuous lifecycle obligation in acquisition policy, not a compliance checkpoint at authorization. This means requiring resilient outcome specifications as a standard element of program management: what must this system do when compromised, in what degraded-mode configuration, and within what recovery timeline? These specifications should be driven by operational commanders and reflected in program requirements, contracts, and sustainment agreements throughout the lifecycle. The acceptable exposure window for vulnerability response—the time between disclosure and effective remediation across the full system lifecycle—should be specified, measured, and enforced as a program performance metric rather than left as an implicit function of program office capacity.

Mandate ecosystem-level threat modeling as a standard element of program protection planning. A program protection plan that does not model the full program ecosystem—contractor organizations, supply chain dependencies, enterprise infrastructure underlying development, workforce dependencies, and the seams between them—is incomplete. Recognize panarchy dynamics in acquisition reform: policy changes at the slow large cycle must deliberately enable the organizational medium cycle to adapt, or technical-layer advances will continue to be constrained by the organizational geometry that produced them. Fund development of scalable resilience guidance that operates at the sociotechnical and ecosystem levels. The technical layer is well-served. The others are not.

For Acquisition Professionals and Program Managers

Expand your aperture from system-centric to ecosystem-level threat modeling. Map the flows and seams in your program ecosystem: between your office and your prime, between your prime and second-tier suppliers, between your enterprise infrastructure and your development pipeline, between legacy systems being integrated and new platforms, between your development organization and your sustainment enterprise. Conway's Law means you cannot eliminate these seams—but you can know

where they are, how wide they are, what traverses them, and what must survive across them when they are exploited.

Treat organizational design decisions as security decisions. because they already are, whether or not you recognize them as that. The structure of your program office, the boundaries of your contracts, the division of responsibilities between development and sustainment: each is simultaneously an organizational choice and an architectural choice with lifecycle security consequences. Make those choices explicitly, with resilience in mind. Ask your operational customer what the system must do when it is compromised. Ask what degraded-mode operation must look like at each stage of the lifecycle. Then engineer toward those answers from program inception, not at authorization.

For Operational Commanders

Assert operational ownership over resilience requirements across the lifecycle. The acquisition system will not produce resilient outcome specifications until commanders ask for them—explicitly, in requirements documents, in contract terms, in sustainment agreements, and in the program reviews where those requirements are tested. The question "What must this system do when compromised, in what degraded-mode configuration, within what recovery timeline?" is an operational requirements question before it is an acquisition question. It belongs to the commander before it belongs to the program manager, and the operational commander reserves the right to change it over the course of the lifecycle in response to adversary adaptation.

The cyber forces that matured over the past decade are among the genuine achievements of the last ten years. They provide detection, response, and deterrent capabilities the joint force did not have before. But their effectiveness at the technical layer is bounded by the resilience of the sociotechnical and ecosystem layers they defend. A cyber force that can detect a compromise faster than the acquisition system can remediate it is a fire brigade in a building with no fire-resistant construction. Commanders who understand both layers—and who drive requirements at both—are the demand signal the next decade requires.

DESIGNING FOR THE INEVITABLE

The decade that began with "identify vulnerabilities and close them" produced exactly what it set out to produce. The defense cyber enterprise has mature technical-layer capabilities, operational cyber forces, and frameworks that have genuinely raised the floor. That is not a small accomplishment, and it is the foundation from which the next decade must build—not a baseline to be criticized.

What the past decade's work has made visible is the layer beneath it. The sociotechnical and ecosystem dimensions of cyberphysical system resilience—the dimensions where Conway's Law operates, where panarchic dynamics determine whether fast-cycle disruptions cascade or are absorbed, where the lifecycle seams from development to sustainment are either managed deliberately or left for adversaries to exploit—these are not exotic research problems. They are the consequential problems that a technically mature defense cyber enterprise is now positioned to address with the same deliberateness it brought to the technical layer a decade ago.

The insight that the attack surface is inherent is not counsel for despair. It is an invitation to ask a more productive question. If organizational complexity inevitably produces seams, and panarchic

cross-scale dynamics ensure those seams can be exploited from the fast cycle to the slow, then the acquisition community's job is not to close all the seams. It is to specify what must survive across them, build organizations capable of adapting when they fail, and maintain a response tempo that keeps pace with adversary exploitation across the full lifecycle of cyberphysical systems. These are achievable objectives. They require policy, doctrine, and sustained investment that have not yet been directed toward them with the urgency the threat environment demands.

The same logic applies to the opportunity AI presents. It amplifies the consequences of Conway's Law and raises the stakes of cross-scale ecosystem vulnerability. It also offers, for the first time, analytical tools adequate to the scale of the problem. An acquisition community that invests in ecosystem modeling, resilience simulation, and supply chain visibility will have a materially better understanding of its attack surface than one that does not. The question is whether the policy and organizational conditions exist to deploy those tools—or whether the same organizational seams that create the attack surface will also prevent us from seeing it clearly.

The attack surface is us. It always has been. The question is no longer whether we recognize it. *The question is whether we design intentionally for it—across the full lifecycle, at every scale, before the adversary forces the reckoning.*

REFERENCES

Allen, Craig R., David G. Angeler, Ahjond S. Garmestani, Lance H. Gunderson, and C. S. Holling. 2014. "Panarchy: Theory and Application." *Ecosystems* 17 (4): 578–589. https://doi.org/10.1007/s10021-013-9744-2.

Atlassian. 2022. "Confluence Security Advisory 2022-06-02," June 2, 2022. https://confluence.atlassian.com/doc/confluence-security-advisory-2022-06-02-1130377146.html.

Cherns, Albert. 1976. "The Principles of Sociotechnical Design." *Human Relations* 29 (8): 783–792.

Conway, Melvin E. 1968. "How Do Committees Invent?" *Datamation*, 28–31.

DoD (U.S. Department of Defense). 2022. *Summary of the Joint All-Domain Command & Control (JADC2) Strategy.* https://media.defense.gov/2022/Mar/17/2002958406/-1/-1/1/SUMMARY-OF-THE-JOINT-ALL-DOMAIN-COMMAND-AND-CONTROL-STRATEGY.PDF.

Hilger, Ryan. 2025. "An Algorithmic Semantic Analysis of Cyber Security and Resilience Guidance against Interdisciplinary Understanding of Resilience Concepts across Time and Scale." PhD diss., Colorado State University. https://www.proquest.com/openview/3fb74588ca759b4a0b22ff4e078f3f42/1.

Holling, C. S. 2001. "Understanding the Complexity of Economic, Ecological, and Social Systems." *Ecosystems* 4 (5): 390–405. https://doi.org/10.1007/s10021-001-0101-5.

Kott, Alexander, and Igor Linkov, eds. 2019. *Cyber Resilience of Systems and Networks.* Springer International Publishing. https://doi.org/10.1007/978-3-319-77492-3.

Mitchell, Steven W. 2014. "Transitioning the SWFTS Program Combat System Product Family from Traditional Document-Centric to Model-Based Systems Engineering." *Systems Engineering* 17 (3): 313–329. https://doi.org/10.1002/sys.21271.

O'Neill, Patrick Howell. 2022. "How China Built a One-of-a-Kind Cyber-Espionage Behemoth to Last," February 28, 2022. https://www.technologyreview.com/2022/02/28/1046575/how-china-built-a-one-of-a-kind-cyber-espionage-behemoth-to-last/.

Trevithick, Joseph, and Tyler Rogoway. 2021. "F-22 And F-35 Datalinks Finally Talk Freely With Each Other Thanks To A U-2 Flying Translator," April 30, 2021. https://www.twz.com/40380/f-22-and-f-35-datalinks-finally-talk-freely-with-each-other-thanks-to-a-u-2-flying-translator.

USD(A&S) (Office of the Undersecretary of Defense for Acquisition and Sustainment). 1987. *Report of the Defense Science Board Task Force on Military Software.* https://apps.dtic.mil/sti/tr/pdf/ADA188561.pdf.

USD(R&E) (Office of the Undersecretary of Defense for Research and Engineering). 2018. *Design and Acquisition of Software for Defense Systems.* U.S. Department of Defense. https://apps.dtic.mil/sti/pdfs/AD1048883.pdf.

WEST Conference. 2019. "WEST 2019: 13 February Luncheon Keynote Address: The Honorable Thomas B. Modly." https://www.youtube.com/watch?v=1P7XA-7AsMc.