Cyber Resilience for Navy Tactical Platforms

Hilary L. Hershey, Camille R. Daniel, and James D. Miller

ABSTRACT

Most modern U.S. Navy platforms could lose critical warfighting capabilities as a result of failure of computing systems, networks, or automated control systems, collectively referred to as "cyber" systems. A failure of cyber systems that control physical ones could cause equipment or vessel damage, or endanger the crew. This article discusses design and operational aspects of adding resilience to modern warships and aircraft that are incontrovertibly dependent on cyber systems.

INTRODUCTION

On 18 August 1943, USS *Abner Read* (DD 526) struck a mine while hunting Japanese submarines off the Aleutian Island of Kiska (see Fig. 1). The explosion caused severe damage, severing the stern, which hung by the starboard shaft for a few minutes before finally sinking to the bottom of the ocean. The rest of the ship survived and was towed into Adak for temporary repairs. That's resilience. A ship with a hole in it literally the size of a cross-section of the ship remained afloat, protected its crew from the elements until it could be towed in for "repairs," and rejoined the war effort 4 months later, providing crew training and then participating in the Battle of Leyte Gulf.[1]

Consider the design requirements for USS *Abner Read*. There were speed, capacity, endurance, firepower, and habitability requirements but probably no specific resilience requirement—certainly not one that said, "The ship must remain afloat if it loses the aft 77 ft. of a 217-ft. hull." Instead there was a recognition that the hull would be subjected to unspecifiable calamities and must incorporate design philosophies that would increase the likelihood of survival and mission accomplishment. Some of these include watertight compartmentation, redundant power generation and distribution systems, comprehensive damage control systems, and robust structural architecture. Similarly, the crew knew their ship and had trained extensively on damage control and recovery from unexpected casualties.

These resilient engineering practices still exist and continue to be refined, as evidenced by more recent catastrophes such as the USS *Cole* attack, the USS *San Francisco* grounding, or the Southwest Airlines Flight 1380 engine explosion.

Enter cyber—or, more correctly, the significant and pervasive reliance on information processing, networks, and data management to maintain, operate, and fight on naval ships, submarines, and aircraft. Systems on USS *Abner Read* were engineered on the basis of more than a century of civilian and military experience with steam ship design and half a century of experience with electrical power generation and distribution systems. Arguably, we are entering the third decade of critical reliance on cyber (see, for example, Ref. 2: "On 21 September 1997, while on maneuvers off the coast of Cape Charles, Virginia, a crew member entered a zero into a database field causing an attempted division by zero in the ship's Remote

Figure 1. General view of the mine damage to USS *Abner Read* upon arrival to Adak on 20 August 1943. (U.S. Navy photo.)

Data Base Manager, resulting in a buffer overflow which brought down all the machines on the network, causing the ship's propulsion system to fail."). However, only in the last few years has mainstream acquisition accepted cyber resilience as a specific objective. Rather than wait another 20–70 years for cyber resilience to be fully incorporated into requirements, we suggest jump-starting this process by translating well-understood resilience concepts into cyber concepts, eventually developing a systematic understanding of cyber that pervades ship and system design in the same way that understanding of mechanics and power does. When considering the cyber adversary landscape, imagine, much like you would when considering traditional adversaries, that you know something very bad will happen 2 years after a given Navy tactical platform is entered into service. What will the U.S. Navy, acquisition program office, and system architects, designers, and developers wish they had done differently today? A premortem approach[3] asks what can be done today to prioritize the following:

1. Could the platform get back to port?
2. Could the platform still fight?
3. Could the platform be restored to full mission capability?

In prioritizing the three questions above, a strong underlying input lies in the cyber resilience of the platform itself.

Another critical attribute of cyber resilience is that it takes advanced persistent threats (APTs) into consideration. An APT is an attack in which an unauthorized user gains access to a system or network and remains there for an extended period of time without being detected. APTs rarely cause damage to networks or local machines; instead, the goal is most often data theft. APTs typically have several phases, including hacking the network, avoiding detection, constructing a plan of attack, and mapping network/system data to determine where the desired data are most accessible. Then, these sensitive data are gathered and exfiltrated. However, with a cyber-resilient design, resilient hygiene procedures in place, and an overall resilient architecture, responsive awareness, cyber resilience, and pervasive agility can thwart the APT adversary.

DESIGN FOR RESILIENCE

Most modern U.S. Navy platforms could lose critical warfighting capabilities as a result of cyber system failure. A failure of cyber systems that control physical ones could cause equipment or vessel damage, or endanger the crew. When a platform's location or mission makes it unreasonable or undesirable to conduct complete, methodical troubleshooting and repair, designed-in resilience must ensure opportunities for minimum recovery that can be executed quickly and locally so that the platform can return to the mission (perhaps in a degraded capacity), be repaired more completely, or as a last resort, limp back to port.

Figure 2 illustrates the resilience of the tactical platform's systems during a cyberattack. Prior to t_0, the system(s) are executing at full functionality. At t_0, the cyberattack occurs, degrading the system's functionality considerably. A cyber-resilient system can reconstitute its critical (threshold) functionality and operate in a degraded capacity during the time period from t_0 to t_1 until full capability can be restored. The goal is to minimize the length of elapsed time between t_0 and t_1.

Because the landscape of cyberattackers is ever evolving and essentially impossible to forestall using current Navy acquisition processes, it is very likely that a successful cyberattack will occur on a tactical platform. Given this high likelihood, our Navy must design tactical platforms to have a highly resilient cyber posture. Therefore, all systems that rely on processing, networks, and data management for full capability should be able

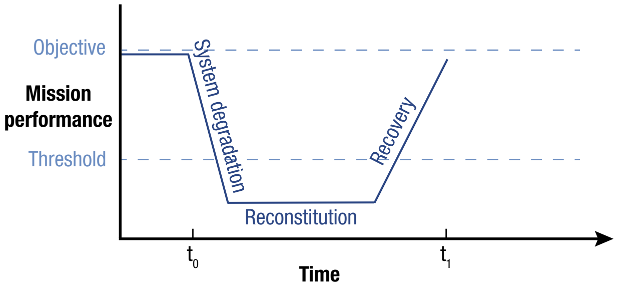

Figure 2. Resilience during a cyberattack.

to withstand attacks and reconstitute functions from an attack-induced casualty. Systems must:

- Withstand attacks
- Detect and alert operators early when there are system and software casualties
- Retain critical (safety and self-defense) functions, preferably with the most minimal reliance on cyber systems possible
- Decouple functionality to reduce or prevent the spread of failures
- Permit rapid recovery at sea
- Evolve to changes in the technical, operational, or threat environments

Withstand Attacks

We offer the following definition of cyberattack: an attack, via cyberspace, targeting an information system's use of cyberspace for the purpose of disrupting, disabling, destroying, or maliciously controlling a computing environment/infrastructure, or destroying the integrity of the data or stealing controlled information.[4] An important nuance is that an attack is an attempt, regardless of whether it has an effect on the system. It is analogous to the legal definition of assault, which occurs when the punch is attempted and is unrelated to whether the punch lands (battery).

Examples of features that enhance a system's ability to withstand a cyberattack include the following:

- Hardware and/or software boundary defenses
- Limited access and visibility to control systems and tactical networks
- Controls to prevent the running of unauthorized applications
- Hardware/software authentication
- Process solutions like protecting system design documentation and installation software
- Periodic rebooting as a precautionary measure against undetected malware
- Architecture, hardware, and software hardening

Detect and Alert Operators Early

While an operator will certainly recognize total system failure, improvements in technology, system knowledge, training, and operating culture are necessary for an operator to reliably detect sophisticated or subtle attacks. Some concepts to consider include the following:

- Early detection of malware and other anomalous behaviors
- Routine comparison to "out-of-band" indications—for example, does the digital course agree with the magnetic course, corrected for variation and deviation for purists (+ deviation)? Do sonar contacts move across the display appropriately for course changes?
- Independent read-back of presets and input parameters

Retain Critical Functions

Designing a system to retain critical functions offers real opportunity. This is far less cyber and more systems engineering. Unfortunately, there is little agreement, few standards, and fewer requirements as to what is critical, how robust the protected capability must be, and how much to invest in protecting these functions. A reasonable standard might be to retain sufficient sensors, indications, and control to maneuver the ship without increasing its exposure to enemy fire.

A recommended methodology is to identify and prioritize the platform's mission-critical functions as indicated below, with the highest priority given to the functionality that may impact ship safety and self-protection (SS/SP) and the corresponding mission-critical functions.

- Ship safety
- Self-protect
- Mission-specific functions
- Convenience functions

To be resilient in the presence of a cyberattack, systems identified as critical to SS/SP functions should have redundant functionality, be capable of being isolated from the tactical network, and be designed to ensure operational capability in a cyber casualty state, preferably with the most minimal reliance on cyber systems possible. There should be no single point of failure for any SS/SP function. Additionally, SS/SP functions should be highly reliable and available even when the rest of the main tactical system or subsystem is down or degraded during a cyberattack. In this casualty state, communications between the subsystems should be through well-defined, specified interfaces.

Decouple Functionality

Leveraging design concepts like modularity, weak links, and redundancy in the design, with the objective to avoid failure propagation,[5] gives the system the ability to fight through attacks. These principles help build the foundation for a more resilient posture for the platform. System architecture and design should limit the extent to which an attack on or failure of one system or subsystem affects other subsystems. Some example concepts include the following:

- Functional independence so failures are not transmitted across interfaces
- The ability to recover or restore subsystems without affecting the rest of the system
- Separate versions of critical ship safety functions to support rapid restoration of these functions

To be more specific, the concept of modularity entails the idea that subsystems should be composed of localized nested modules specific to subsystem functionality. Additionally, each subsystem is self-contained, and communications beyond the subsystem occur only through well-defined, specified interfaces (e.g., AMQP or in legacy systems, CORBA). These two attributes of modularity improve the likelihood that damage to a subsystem can be localized and contained during a cyberattack.

Weak links means that the linkages between subsystems and functional modules will break in a manner that minimizes the chances that a cyberattack originating in one subsystem or functional module can propagate to an adjacent subsystem or functional module.

Rapid Recovery at Sea

The need to recover rapidly at sea is the least unique to cyberattack and the most amenable to requirements. Information processing systems should be recoverable at sea in a tactically acceptable time frame just like any other systems. Physical at-sea recovery requirements should exist for all computing hardware, application software, operating systems, basic input/output systems, and other firmware. Examples include quick restart/node reset, quick reboot/subsystem restart, full reboot, and reset to "gold." In addition to the reset/reboot recovery options, onboard hardware spares should be available in the event of normal hardware failure or a cyberattack that breaks the hardware. Recoverable system components should include all computing platforms (rack mount, laptop, tablet, etc.), networking hardware, disk drives and other storage devices/media, and any other components that could be affected by cyber effects.

Evolve to Changes

Because the cyber threat is ever evolving, it is important to continually evolve the Navy tactical platforms' resilience, strategy, and defenses. By adapting mission functions and supporting resilience capabilities to predicted and existing changes in the technical, operational, or threat environments, tactical platforms are better positioned to address the emerging cyber threat.

TRAIN FOR RESILIENCE

Our experience is that most cyber training is focused on cyber hygiene (i.e., password construction, trusted media, authorized software, etc.) and lacks strong emphasis on casualty response and damage control. At-sea operators require tools, knowledge, and practice to rapidly and effectively address ship casualties. (These practices apply to any casualty, whether it is cyber induced or not.)

A first step toward improving this condition is the publishing of ship casualty procedures for cyberattack in the same publications and instructions that include ship casualty procedures for fire, flooding, loss of power, etc. Doing so brings cyber-related casualties into the mainstream because these publications are the basis for all shipboard training. Procedures should also mirror the form and objectives of the traditional casualty procedures:

- Place the ship/system in a safe condition
- Warn affected personnel
- Isolate affected systems from other systems
- Arrest the spread of the casualty
- Place the affected system in a condition for recovery

At-sea operators need situational awareness and casualty recovery tools that do not require them to be cyber experts. The recovery tools may be as simple as read-only drives to reboot and gold disks to reload, but these should be designed to assist the sailor in recovering as quickly as possible. Some simple suggestions include that reboot/reload be as segmented as the system so that unaffected parts of the system are not lost when affected parts are reloaded.

More challenging is casualty detection. The "blue screen of death" is certainly unmistakable, but more subtle attacks may disable capabilities or cause insidious failures ultimately resulting in injuries or physical damage if operation continues.

For example, false attitude indications may result in some physical limitation being exceeded. To prevent this, the operator may need to routinely compare independent indications (assuming such indications are available), at least one of which preferably does not rely on cyber (e.g., an analog depth gauge for a submarine). Such comparisons should be an integral part of sailing the vessel (much like an instrument scan is a required part of piloting an aircraft in clouds), both in steady-state conditions or to validate responses to control inputs.

Finally, practice. Casualty drills should include cyber-induced casualties. This may require designing in the ability to simulate casualty conditions while normal processing carries on in the background. There is no substitute for training and practice. Cyber-induced casualties must be part of standard certifications and inspections, which should be reviewed at the force level to assess crew performance and to generate recommendations for improved doctrine and procedures.

A different style of casualty drill would be cyber stress testing of both the system itself and the operators to assess the current state of resilience. Netflix began testing its distributed systems in a fashion that deliberately reduces the amount of resources available for streaming—the reduction is from individual servers to entire regional distribution centers. Netflix used the term *chaos monkey*,[6] referring to the idea that one can release a monkey that stresses a system to a controlled point and then cage the monkey again, therefore returning the system to its normal state. Extending the definition of stress beyond the Netflix usage can allow for emulation of certain cyber threats and even performance of some cyber "red team" types of attacks. A cyberattack commonly reduces confidentiality in, integrity of, and/or availability of a system. Accurate modeling and knowledge of how a system works can enable creation of a derivative of the chaos monkey concept that specifically targets a system and the operators' response to cyber attacks.

As in other aspects of Navy tactical platforms, there is a concept of readiness to perform the mission, in this case cyber readiness.

CONCLUSION

Digital processing has greatly enhanced the capability of ship and aircraft systems and in many cases reduced the estimated total cost of ownership. Designing robust digital processing systems has added new opportunities for providing our warfighters with the most capable systems on the planet. In a modern Homeric tragedy, our unbeatable systems have an "Achilles' heel"—the susceptibility to sophisticated and subtle, remote, targeted attacks that can create potentially devastating effects.

To mitigate these potentially harmful effects, the design and architecture of the Navy tactical platform should account for the possibility that an adversary may gain a persistent presence on one or more subsystems, or their components, resulting in degraded mission functions. Cyber resilience enables the platform to continue mission functions, even if in a degraded capacity, until that system's/subsystem's mission functions can be restored to a fully operational state.

Most of the design considerations discussed in this article are best implemented in the design phase. However, because most U.S. tactical platforms were designed decades ago, changes must be backfit. For existing systems, a careful review and upgrade is necessary and should be conducted as prioritized here.

- Ship safety
- Self-defense
- Mobility
- Mission-specific functions
- Convenience functions

If informed by complete understanding of the existing systems, training and operational changes could make huge, quick improvements. The Naval Sea Systems Command (NAVSEA) and Naval Air Systems Command (NAVAIR) should conduct design reviews and develop alterations to provide manual backup for digitally controlled ship systems, with maneuvering, navigation, power generation and distribution, cooling, and freshwater production systems given priority. These systems need to ensure that a ship could return to port if necessary. Second priority should go to damage control and force protection systems. Hardware and software changes will likely be required to add a measure of cyber resilience to more complex systems such as main battery fire control, air and missile defense, and long-range sensors, and these changes should be prioritized with other capability upgrades.

At the same time, casualty procedures for cyber-induced casualties should be a Fleet priority. In many cases of digital control systems, these casualties are versions of traditional casualties caused by attack on/failure of the embedded digital processing.

USS *Abner Read* was not specifically designed to operate or even float with a quarter of the hull missing. It was purposefully designed to withstand and recover from unexpected attacks of a predictable nature. We are suggesting that the ubiquitous digital systems onboard our ships and aircraft warrant the same prudence. The specifics of a given cyberattack are unknowable in advance, but the effects can be generalized, and sensible preparation through training, practice, and design will limit the ultimate consequences, preserving capability, investment, and lives.

REFERENCES

[1] Fuentes, G., "75 Years Later, Search Finds Ship's Stern Ripped Away by Mine in WWII's Aleutians Campaign," *USNI News* (accessed 16 Aug 2018).

[2] Wired Staff, "Sunk by Windows NT," *Wired*, July 24, 1998, https://www.wired.com/1998/07/sunk-by-windows-nt/.

[3] Klein, G., *Streetlights and Shadows: Searching for the Keys to Adaptive Decision Making*, MIT Press, Cambridge, MA, pp. 63, 235–236 (2011).

[4] Kissel, R. (ed.), *NISTIR 7298: Glossary of Key Information Security Terms*, Rev. 2, National Institute of Standards and Technology, Gaithersburg, MD (2013).

[5] Hole, K. J., *Anti-fragile ICT Systems*, Springer, Bergen, Norway, pp. 35–39 (2016).

[6] Brodkin, J., "Netflix Attacks Own Network with 'Chaos Monkey' – and Now You Can Too," *Ars Technica*, https://arstechnica.com/information-technology/2012/07/netflix-attacks-own-network-with-chaos-monkey-and-now-you-can-too/ (accessed 22 Aug 2018).

Hilary L. Hershey, Force Projection Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Hilary Hershey is the technical lead for defensive cyber in the Sea Control Mission Area in APL's Force Projection Sector. She earned a B.S. in computer science and statistics at American University and an M.S. in computer science at the Whiting School of Engineering at Johns Hopkins University. Since joining APL in 1993 after 5 years as a resident subcontractor, she has applied her software systems engineering skills and expertise to solve problems in anti-submarine warfare; sonar and combat system development for submarine, surface, air, and surveillance platforms; submarine operations; periscope detection radar; and applied cybersecurity and resilience for Navy tactical platforms. Her e-mail address is hilary.hershey@jhuapl.edu.

Camille R. Daniel, Force Projection Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Camille Daniel is the manager of the Defensive Cyber Program in the Sea Control Mission Area of APL's Force Projection Sector. She earned a B.S. in mathematical sciences from Spelman College, an M.Ed. in mathematical sciences from Virginia State University, an M.S. in mathematics and applied mathematics from Virginia Polytechnic Institute and State University, and a D.Eng. from George Washington University. Camille joined APL in 2005 and has worked in the areas of anti-submarine warfare analysis, assessments, and requirements; counter ISR assessments and technology development; and defensive cyber testing, assessments, operational impacts, and requirements. Her e-mail address is camille.daniel@jhuapl.edu.

James D. Miller, Force Projection Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

James D. Miller is the manager of the Submarine Warfare Program Area in the Sea Control Mission Area of APL's Force Projection Sector. He earned a B.S. in Electrical engineering from the United States Naval Academy and an M.S. in electrical engineering from Johns Hopkins University. He served 25 years as a submarine officer in the U.S. Navy. Since joining APL in 2004, he has worked in the areas of submarine technology and operations. His e-mail address is james.d.miller@jhuapl.edu.