

Resiliency in Future Cyber Combat

Col William D. Bryant, USAF

The winds may fell the massive oak, but bamboo, bent even to the ground, will spring upright after the passage of the storm.

—Japanese proverb

Abstract

Rigid cyberspace defenses are proving unable to meet advanced and modern cyberspace threats. As a result, there has been increasing focus and interest in cyber resiliency, but what will it take to be resilient in future cyber combat? We can glean some useful concepts from the ancient Japanese proverb about the resiliency of bamboo in a storm. In comparison with the massive oak, which relies on structural strength, three characteristics enable the bamboo's greater resiliency. Bamboo has the ability to accept deformation without failure and a significantly reduced attack surface, and it dynamically reacts to the wind in a way that minimizes the impact of future gusts. Defenders of cyberspace should look to add similar characteristics to their cyberspace systems. First, cyberspace defenders should maximize the flexibility of their systems by deliberately building in "inefficient" excess capacity, planning for and expecting failure, and creating personnel flexibility through training and exercises. Second, defenders should reduce their attack surface by eliminating unnecessary capability in both hardware and software, resist users' desire for continual rapid improvements in capability without adequate security testing, and segment their networks and systems into separate defended enclaves. Finally, cyber defenders should position themselves to dynamically respond to attacks through improved situational awareness, effective cyberspace command and control, and

Col William D. Bryant is a career fighter pilot and strategist with a PhD in military strategy from the School of Advanced Air and Space Studies. He has served in numerous operational and staff assignments and is currently the deputy director of Task Force Cyber Secure on the Air Staff. His recently published book is titled *International Conflict and Cyberspace Superiority: Theory and Practice* (New York: Routledge, 2015).

active defenses. Combining these approaches will enable the defenders of cyberspace systems to weather cyberspace attacks and spring upright after the passage of the storm.



According to the ancient Japanese proverb, after the storm passes, the stronger oak lies on the ground while the weaker bamboo stands upright. The moral that resiliency is more important to success than strength applies to conflict in the cyberspace domain as well. It is important to clarify that the resilience being discussed here is in response to cyberspace attacks, not cyberspace espionage. Cyberspace attacks change friendly systems through manipulating data, causing hardware failures, or physically destroying objects controlled from cyberspace. If pure cyberspace espionage is done well, the defenders will have no idea anyone was ever in their systems: everything will still function. Resilience is not as useful in examining cyberspace espionage as it is in investigating cyberspace attack.

The Department of Homeland Security's Risk Steering Committee has defined resiliency as the "ability to adapt to changing conditions and prepare for, withstand, and rapidly recover from disruption."¹ As organization after organization and system after system is successfully attacked, there is a growing realization that a perfect perimeter defense is not possible, and even if it were, attackers are often within the walls as insider threats. In addition, while shifting to multiple layers of "defense in depth" improves security, each layer will still have flaws and vulnerabilities that a determined attacker can circumvent. Accordingly, cyberspace operators have increasingly looked to resilience as a promising way to improve overall security.²

While resilience is the key to success for cyberspace defenders, it is important that defenders not neglect traditional network defenses. In the US military, the tendency has been to pour a disproportionate amount of resources into offense while not focusing enough on defense. This is a mistake, and as noted by Martin Libicki, "In this medium, the best defense is not necessarily a good offense; it is usually a good defense."³ Offense is widely seen as overwhelmingly powerful over defense, but that assumption ignores the historical record of cyberspace attacks to date. Modest defenses easily defeat unsophisticated attacks, and even nation-state-level attacks have had mixed success. Of the eight cases of nation-

state on nation-state cyberspace attacks with a reasonable amount of open-source data, only half can be qualified as a success.⁴ If the offense were truly so overwhelming, it should be able to achieve greater than a 50 percent success rate. When the high-level attacks are analyzed, it is apparent that in most cases the offenders did get past the defenses, but the defenders were able to react and negate the attacks in a week or two. Resilience is the key to an effective response.⁵

Before developing the tenets of cyberspace resiliency, it is important to clarify what cyberspace is, as there is great confusion on this point. The US Joint Staff has defined cyberspace as “a global domain within the information environment consisting of the interdependent network of information technology infrastructures and resident data, including the Internet, telecommunications networks, computer systems, and embedded processors and controllers.”⁶

One important point emerging from the definition is that while the Internet is part of cyberspace, it is not all of cyberspace. Any computer processor capable of communicating with a computer system is in some way part of cyberspace. A desktop computer, an avionics computer on an aircraft, an iPhone, an industrial controller, and the central processor on a modern car are all part of cyberspace, although only some of them are routinely connected to the Internet. Most modern military equipment is more complex than an M-4 carbine and has some form of processor, from a humble truck to an aircraft carrier, and is thus part of cyberspace. So what is required to be resilient within cyberspace?

Using the bamboo analogy, there are three elements of success against the storm that have application to resiliency in the cyberspace domain: flexibility, a reduced attack surface, and the ability to respond dynamically to attacks. First, the bamboo can accept deformation without failure. As noted by the proverb, the bamboo can be bent and spring back upright, while the oak can accept little deformation before failing catastrophically. Second, the bamboo presents far less attack surface to the attacking wind, as it has a streamlined shape with relatively few exposed leaves compared to the oak tree, which has a far larger and more complex structure. Finally, the bamboo adjusts to the wind, bending to minimize the effect of future wind gusts. Each of these three characteristics can be applied to the cyberspace domain as a way of understanding how practical cyberspace resilience can be achieved.

Flexibility

What does flexibility look like in cyberspace? Is flexibility even a meaningful concept when every device in cyberspace is actually running a complex rule-set that predetermines its actions in response to a given set of inputs? While the computers that make up cyberspace simply do what they are told, the flexibility in cyberspace comes from people telling machines what to do. People can also build in more capacity for flexibility by constructing their systems to operate in cyberspace with excess capability.

The typical business mind-set focuses on efficiency to generate as much profit as possible, while the military mind-set loves both efficiency and order, but both concepts are antithetical to flexibility in cyberspace. Efficiency means using 100 percent of available resources with no excess capability. Yet, if you are 100 percent efficient, the smallest perturbation can lead to catastrophic failure. The heart of resiliency is the ability to absorb perturbations and failures—whether natural or manmade—and continue functioning. Thus, a system built for resiliency will look very different than one built for efficiency.

Too much efficiency will hamper resiliency, and cyberspace defenders would do well to build less-efficient redundant systems if they want to achieve resiliency under attack. Cyberspace operators who want to build a resilient system must oppose several efficiency trends. In the perfectly efficient network, every device on the network will run the same operating system, utilize the same applications for specific tasks, have a minimum of subnetworks or enclaves all structured the same way, and even utilize the same hardware throughout for the same functions. While these concepts are efficient, they are not resilient.

An entire network that runs a single operating system is efficient and easy to administer and also just as easy to take down via a single vulnerability. A heterogeneous network made up of Windows 7, Windows 8, and Linux—with a few Apple machines thrown in for good measure—cannot be completely taken down by a single vulnerability. Military strategist Edward Luttwak noted that in the strategic realm with a thinking enemy “homogeneity can easily become a potential vulnerability.”⁷ Of course, there must be a balance between efficiency and resiliency; a system where every single device runs a unique operating system would be resilient in a sense but would be so difficult to administer it would likely be full of unpatched machines and unknown

vulnerabilities. Aristotle taught that virtue always lies on a continuum between two vices.⁸ The virtue of cyberspace resilience lies between rigid conformity to a single system that can be taken down with a single attack on one side and complete chaos within an unworkable mess of a network on the other. A reasonable middle ground for cyberspace operators is to select a handful of different, well-designed operating systems and then implement them throughout their networks. Thus, if three operating systems are used, two-thirds of the network should be available following any attack that uses a single vulnerability. Cyberspace operators should also find the right balance between too many and too few different types when it comes to applications and hardware, for very much the same reasons as discussed above for operating systems. Heterogeneous systems are a start, but the defender can also break those systems into separate enclaves to further increase resiliency.

Network segmentation into separate subnetworks that can function even if other networks around them fail is a key component of cyberspace resiliency. The current push toward ever-larger homogenous networks is good for efficiency but not for resiliency. Consider the changes the network on a typical military base has gone through. At first, every base was unique; information technology (IT) equipment was purchased locally, and every network ran different software and applications depending on what the local communications unit purchased. This structure was extremely inefficient, and the level of security achieved was highly variable and often quite low—partially because the different networks still had to be connected to each other, often in not very secure ways. The next step was to bring control of the networks up to a regional level, which took control of base networks largely out of the hands of local units. While this resulted in a more efficient network, it also meant a successful attack against the regional hub could bring down multiple bases at the same time, whereas before, each base would have to have been reconnoitered and attacked separately. Now, the Department of Defense (DOD) has mandated that the military services all utilize the same structure under the Joint Information Enterprise and the same network hubs in the form of Joint Regional Security Stacks. If every service is using the same equipment running the same software, one successful attack against a single vulnerability could conceivably take down the entire DOD network.

Returning to the “Wild Wild West” mind-set, where every local unit does whatever it wants, will not improve operational effectiveness. Instead, resiliency and a reasonable level of efficiency can be achieved by a deliberate diversification of networked systems. Homogeneity is good for ensuring patches and protocols are followed against known threats, while heterogeneity helps protect a system from unknown and unpredictable threats. System architects should buck the trend toward ever-larger and homogenous networks and deliberately build in heterogeneous enclaves based on a small number of carefully selected hardware and software configurations. It is important that network architects do not build systems with a single type of system performing a function across the network. For example, a segmented network of heterogeneous enclaves that all use the exact same hardware and software as a gateway will be less resilient than one that uses different types of gateways. Resiliency is best increased by parallel lanes of different systems, if a network relies on a single type of system at any level, there is still a single point of failure. As with operating systems, finding Aristotle’s “golden mean” of enough diversification to be resilient with enough efficiency to be manageable and low-cost is the key.

Even if a network is heterogeneous and cannot be completely taken down by a single vulnerability, cyberspace operators still need to expect and plan for failure.⁹ Planning for failure does not come naturally, especially in the military environment. Complexity theorist Antoine Bousquet has noted that the military often attempts to achieve “100% relevant content, 100% accuracy, and zero time delay” in the pursuit of a frictionless cybernetic war machine, but that goal is illusory. Instead, cyberspace operators should be “embracing uncertainty and designing a resilient and flexible military that is capable of adapting to the unforeseen and contingent.”¹⁰ Cyberspace operators need to move beyond the concern of how to best secure their systems against attack to focus on how to design their system to continue working after their defenses fail. This requires a significant mind-set shift for military cyberspace operators, including focusing on response capabilities such as emergency and incident response teams and plans.¹¹ One of the best ways to accomplish this shift is through aggressive and thorough red teaming.

Well-resourced and extensive red teaming of cyberspace systems is a critical part of building cyberspace resiliency. A red team is a group of friendly attackers who attempt to attack systems to find their vulner-

abilities and weaknesses. They use the same techniques as real attackers and provide an invaluable service in not only finding vulnerabilities but also giving defenders practice in how to respond to attacks and keep their systems functioning. To get the maximum benefit out of red teaming, exercise referees need to allow red teams to breach defenses and actually do damage within the exercise system; stopping the exercise when the red team gets access does not yield as much benefit. Historically the DOD has underresourced red teams due to the persistent focus on offensive cyber capabilities. Red teams require the same people and resources needed for offensive cyberspace capabilities. However, offensive capabilities and red teams are not locked in a zero-sum resource game. Since the same attack techniques are used, red teaming can be excellent training for offensive cyberspace operators and can help overcome classification barriers.

Compartmentalization continues to be a major issue preventing defenders and attackers from learning from each other.¹² According to former vice-chairman of the Joint Chiefs of Staff, Gen James Cartwright, “We make sure the recce teams don’t tell the defenders what they found, or the attacker, and the attackers go out and attack and don’t tell anybody they did. It’s a complete secret to everybody in the loop and it’s dysfunctional.”¹³ Compartmentalization and security are essential in protecting cyberspace weapons, but it is foolish for attackers to assume their enemies will not discover and utilize the clever techniques they develop. Attackers need to inform friendly defenders of their attack methods in appropriate ways that allow defenders to defend their systems, while not giving away the attack methods to adversaries. Once again, there is a balance required between disclosure and security, but it appears in the DOD the needle is too far toward security. More disclosure by attackers to defenders is needed for improved cyberspace resiliency. Red teaming and improved disclosure helps to develop resiliency in the people operating in cyberspace, but there are a number of other ways to build resiliency into cyberspace operators.

The highest payoff in building cyberspace resiliency lies in building resilient people. People, not machines, react. The machines will simply do what their instructions tell them to do, even if those instructions are complex and allow for some ability to respond to stimuli. Not only do cyberspace operators need to be resilient, improved resiliency and security needs to be built into system users as well.

Education that creates deeper understanding of the cyberspace environment will often yield a major payoff in unexpected ways and places. Training is valuable and can produce an immediate payoff; education takes longer but can provide more benefit in the long run. Because of the long-term and difficult-to-measure nature of the payoff, education often takes the first hit when an organization is under budgetary pressure. This is shortsighted and will reduce an organization's abilities to understand the environment and to generate the cultural change needed to build an organization that can be resilient in the cyberspace domain. Education lays the foundation, but training provides the specific tools needed by people operating in cyberspace.

Users can be "hardened" via training, as they are currently the weakest spot in the armor of most cyberspace systems. Users are the bane of system administrators the world over, and many attacks rely on finding a user who can be tricked into compromise. Most users have only a rudimentary knowledge of computer security, so spending time and money training them can produce a significant payoff. Mandatory training programs are a start, but not all users will pay attention to training or be convinced that it is important to them. System administrators need to convince users there is a significant benefit to following good security practices, whether it is monetary rewards for best practices or reprimands for those who do not follow procedures.

Most organizations have a user training program in place; what is missing is accountability to make users take cyberspace security seriously. In a recent study, security testers left USB thumb drives on the ground in a parking lot outside of a federal office building. All federal employees receive regular training on the dangers of plugging in unknown USB devices, but 60 percent of these highly trained employees plugged them in anyway. The addition of an official looking logo on the drive increased the percentage of USB drives employees plugged in to 90 percent.¹⁴ How many of these employees were fired or even mildly reprimanded for their failure to follow procedures? Performance ratings and rewards need to be explicitly tied to following security practices, and there should be consequences for security failures that are regularly tested via a continuing testing program.

Users should be routinely tested and probed, and those who do not perform well should face escalating consequences. For example, cyberspace operators should routinely send out "phishing" style e-mails to users

of their systems based on actual real-world attacks. If a user is duped into clicking on the link, instead of unleashing a virus, the user should be directed to retake the organization's computer security training. Subsequent failures should have increasingly unpleasant consequences, including eventual termination for employees who are incapable of following good security practices. A similar escalation ladder could be followed for users who continue to visit questionable sites or are caught deliberately circumventing security safeguards. Escalating consequences are for well-intentioned but security inept employees; insider threats are a different matter and should be dealt with according to organizational and legal rules. These types of changes will normally not be received well because they involve changing organizational culture and they will require support from top executives in the organization to be successful. For the military in particular, cyberspace resilience will also include a significant amount of resilience outside of cyberspace systems.

For most Western militaries, cyberspace systems are principally important because they enable effectiveness in the physical domains of combat. Thus, cyberspace resilience includes the ability of military forces to fight effectively even if their cyberspace systems are compromised or unavailable.¹⁵ Management scientist Martin Libicki has recently identified that networked militaries need to be careful not to focus on the network for its own sake through information assurance but must instead stay focused on the mission and mission assurance.¹⁶ This is deeply uncomfortable for a generation that has become accustomed to continual connection and reliability of cyberspace systems, since Western militaries have not yet fought a significant cyberspace adversary. However, there are a number of potential adversaries who have been very clear that they intend to fight hard in cyberspace in the case of a conflict, and Western militaries would be exceedingly foolish to assume the enemy will never have a "good day" and be able to disrupt many critical systems.

Much like with cyberspace operators and system users, resilience in regular military forces can best be built through realistic training and exercises. Currently, if there is cyberspace play in military exercises, exercise referees usually discount it so regular forces can receive "good training" and utilize all their systems. On the contrary, *good training* is when they do not have all their systems. Consider a single cyberspace enabled system, the Global Positioning System (GPS). What would a major exercise such as Red Flag look like if none of the participants were

allowed to use GPS? What about a land-based combat exercise at the National Training Center? How many young platoon leaders would be able to maneuver their forces quickly and expertly using only a compass and a map? What if their radios stopped working as well? Would forces continue to maneuver and operate in the absence of communication from headquarters? There are some hopeful signs that some leaders in the military are taking this threat seriously, and distributed control is one promising approach.¹⁷ But these ideas must be thoroughly tested and exercised on a regular basis if military forces are going to have any ability to operate in a cyberspace-denied environment. The flexibility created by these changes can be enhanced by also reducing a cyberspace system's attack surface.

Reducing Attack Surfaces

Bamboo has far less surface area in its structure than the massive oak tree and thus presents a much smaller surface for the wind. In cyberspace, the surface area is normally referred to as the "attack surface." The attack surface is made up of all the potential access points for an attack. Cyberspace operators should actively seek to make their attack surface as small as possible so the effect of each attack and the resultant recovery time are minimized. The fewer systems that must be recovered, the more quickly recovery can take place.

Every piece of software, every capability added to that software, and every communications pathway represents a potential avenue of enemy attack. Thus, the first thing cyberspace operators should do to reduce the attack surface they present to the enemy is eliminate nonessential features, as such features represent added risks.¹⁸ This is a mammoth task—one most cyberspace operators are not easily able to accomplish, as modern software is written to appeal to the largest number of customers with all the bells and whistles on by default. It can be nearly impossible, not only to determine what functionality is unneeded but also to disable it across the network to prevent it from being used as an attack vector. While swallowing the elephant all at once is not immediately achievable, concrete steps can be taken in both the software and hardware arenas.

An organization's hardware attack surface can be reduced by disabling unnecessary ports and communications pathways where possible. The best method for disabling unnecessary communications pathways is

via physical means. If a computer's wireless network or modem card is removed, there is complete certainty that card cannot be used as a clandestine back door into the organization's network. The same can be said for unnecessary physical connections such as USB ports. It may be inelegant to cut the wire behind the port, or fill the port opening with hot glue, but it is effective. Software methods can be used as well, even if they are not as effective. Most devices can be easily disabled via the operating system, although cyberspace operators would do well to run periodic checks to ensure disabled ports have not been turned back on surreptitiously. While closing hardware-based vulnerabilities is a start, the majority of an organization's attack surface lies in its software.

If an organization is serious about its cyberspace security, there should be an increased level of scrutiny of every piece of software on the organization's networks. As with operating systems, there must be a balance between having too many and too few different applications to accomplish the same function. Having too many applications opens unnecessary avenues of attack, while having too few can hamper resiliency. For example, if an organization mandates the use of only one particular build of one Internet browser, it is difficult for the organization to react quickly if a vulnerability is discovered in that browser. If the same organization had two browsers on the network, all functionality could be quickly switched to the second browser while the first was patched. However, in many networks there are clearly far too many applications, not too few. It is reasonable to have a primary and a spare application for each function, but not reasonable to have 12 applications that do the same thing. Of course, a system administrator telling users they cannot utilize their favorite application is about as popular as the Internal Revenue Service auditor. The importance of reducing an organization's attack surface presented via too many applications is not well understood. Users will often push back against security requirements if it means they cannot get the software they want as quickly as they want to get it.

There is a natural tension between the desire of users for continual connection with constant improvements and security requirements to restrict unnecessary communications pathways and comprehensively check all new software. Once again, balance is the key, and cyberspace operators must find the correct balance between competing requirements. That balance will be different for each organization and operational environment. The right balance for a small IT firm developing

iPhone applications is very different than the right balance for the National Security Agency (NSA). Once an organization has done all it can to reduce its software and hardware attack surfaces, it can take one more step.

The final step in reducing the attack surface shown to a cyberspace attacker is to hide as much of it as possible behind different segments of the network. In many respects, this is very similar to a “defense in depth,” where once the attackers get over one wall, they are faced with a whole new series of walls different than the first.¹⁹ These additional barriers can also provide more opportunities for the defender to detect the attack. There is some overlap with the previous discussion on flexibility, as segmentation can aid flexibility and also reduce an organization’s apparent attack surface. True, air gaps remain very difficult to maintain without some connection for maintenance or communication, but segmenting a network into different areas with strictly controlled communication links can reduce an organization’s attack surface. The amount of communication allowed into or out of the segments can be adjusted to account for the level of security required. The control system for a nuclear power plant should have very little and strictly controlled access to communication flows, whereas the segmented network for an operating division inside a corporation will normally have much freer communication links.

The key to reducing attack surfaces appropriately is finding the right balance between connectivity and security. However, the world is full of aggressive actors in cyberspace, and it is likely an attacker will find a vulnerable attack surface eventually. Then the flexibility an organization has built will be tested as it reacts to the storm and bends like the bamboo.

Reacting to Attack

When bamboo bends during a storm, it is in response to the pushing of the wind. The bamboo bends away from the wind, which reduces the amount of the bamboo’s surface the wind can push on, and the bamboo’s leaves and branches streamline, which reduces the force on the bamboo even more. If cyberspace operators are going to react to attacks in an analogous way, they must start by understanding what is going on around them in cyberspace.

Cyberspace situational awareness is normally essential for cyberspace operators to effectively react to an attack. Defenders must know they are under attack before they can react resiliently. If an attacker is simply trying to steal information or implant pathways for future attacks, he or she may work very hard to avoid detection. Resilient cyberspace defenders must find the enemy to affect a response, but to find the enemy, defenders first must know their own home terrain.

Cyberspace situational awareness starts with the cyberspace defender; cyberspace operators need to understand their own networks. This point at first may seem so obvious as to be hardly worth stating, but it may surprise people outside the IT industry that many large organizations—including the military services—do not have a complete picture of what their networks look like, exactly how and to what they are connected, or even what applications are running on their networks. Large amounts of money are being spent to attempt to solve this problem, but an immediate solution is not apparent. Automated tools do a good job finding what they know to look for, but unique systems and applications are often missed, as the automated tools do a poor job of finding and categorizing unknown software. Legacy networks can be riddled with “servers” that are actually desktop computers sitting under a desk somewhere that were configured years ago to do a specific task that may, or may not, still be required. Once a picture of your own network is built, the next step is to consider what the enemy may be trying to do to it.

Intelligence on cyberspace threats is extremely difficult to collect. Cyberspace weapons are easy to build in extreme secrecy, as the resources needed to create them can be easily hidden compared to the resources required to build a battleship or bomber. Intelligence agencies should pursue information on cyberspace capabilities and intentions, but much of their best work will likely not be via cyberspace but via other, more traditional methods, since people drive cyberspace. Moreover, people may yield better intelligence than computers and networks in this area. Consider for a moment the presumed difficulty a nation-state would have hacking into the NSA compared to how easy it apparently was for Edward Snowden to walk out with an enormous amount of information. Once intelligence has been collected, via whatever means are available, cyberspace operators must overcome their own organization’s security policies if it is to have any real effect.

While cyberspace weapons are very vulnerable to compromise and must be protected to be successful, the current extreme levels of secrecy hamper cyberspace resilience. Cyberspace capabilities were developed in a world steeped in high levels of classification and compartmentalization, and it is true that cyberspace weapons are very frangible. Once an enemy knows about an exploit or technique, the target can normally block it very quickly.²⁰ However, a better balance between security and strengthening defenses needs to be struck in this area. Cyberspace attackers should not have to share the details of their latest weapons and techniques, but they should provide generalized threat information based on friendly weapons for the benefit of their own cyberspace defenders. While an enemy might not use identical weapons, similar attacks might be thwarted. For cyberspace attackers to assume their enemy could not possibly be smart enough to discover the same vulnerabilities and techniques would be extremely foolish. Senior leaders who can balance improved defenses against possible loss of offensive capability will have to strike the right balance in this area. Sometimes the answer will be to disclose, and sometimes the offensive capability will be so important that the risk to friendly networks of leaving them unpatched will be deemed acceptable. Right now, it appears the default is that offensive forces share very little with their defensive brethren. For most organizations, Aristotle's golden mean appears to lie in the direction of more disclosure, not less. The next quandary for cyberspace operators is how to effectively command and control cyberspace forces.

Resilient operational cyberspace organizations should be commanded and controlled more like maneuver forces in the physical domains than managed as IT departments. Despite protestations by some analysts that there is no maneuver in cyberspace, humans, who make decisions and react to their adversaries in ways that would still be familiar to Carl von Clausewitz and other military thinkers, continue to drive conflict in the cyberspace domain.²¹ Attempting to reduce military conflict to an engineering problem was a bad idea in the physical domains. Why would we expect it to be a good idea in cyberspace? Accordingly, structuring cyberspace forces as maneuver units that are expected to react and maneuver to defeat a thinking and reacting adversary is a good start. Currently, cyberspace command and control is also far too complex, with decisions to employ too far up in the chain of command and examined by too many different teams of lawyers. Streamlining the process is important,

but that will likely take time, as understanding of the cyberspace domain continues to develop. Meanwhile, maneuver and counterattack will remain important tools for resilient defenders.

According to Clausewitz, defenders should not simply wait passively; “the defensive form of war is not a simple shield, but a shield made up of well-directed blows.”²² William Owens, Kenneth Dam, and Herbert Lin demonstrated that with only a passive defense the defenders have to succeed every time, and since there are no penalties for the attacker, he can continue attacking until he is successful. This difference places “a heavy and asymmetric burden on a defensive posture that employs only passive defense.”²³ A defender can be attempting to accomplish several things when counterattacking in cyberspace. A defender can disable the computers executing the attack and attempt to trace an attack back to its source. Attackers will normally bounce attacks through multiple servers to attempt to hide themselves, but a persistent defender can sometimes work back through the servers to the source or use more creative methods to identify the attacker. If a defender makes it to the originator of the attack, there are now a number of unpleasant things he or she could theoretically do to the attacker’s networks in retaliation. Unfortunately, most of those things are currently illegal for defenders to do under US law.

Since active defense normally involves breaking into a number of privately owned computers along the way, it is generally illegal under the Computer Fraud and Abuse Act. According to Paul Rosenzweig, any active defense that reaches outside of the defender’s computer system is “almost certainly a crime in and of itself.”²⁴ This legal issue opens cyberspace defenders up to prosecution and lawsuits, whether they are military or civilian. And that is just if the attacking computers are only in the United States, which is normally not the case. Breaking into computers in foreign countries brings on entirely new sets of legal and political problems. The difficulty in attributing attacks might work in the defender’s favor, as it can be hard to attribute “hack backs” if the defender chooses to mask where he or she is coming from, but that does not actually deal with the legal and ethical issues. Hack backs quickly devolve into a legal and political Gordian knot. Hack backs are a key element in an effective defense, but they are clearly illegal. However, it is just as clear that even private organizations are now using them.²⁵ Hopefully, policy and legal authorities will catch up in this important

area. Fortunately for defenders, there are other types of active defenses that pose fewer legal issues.

A less legally problematic technique a resilient cyberspace defender can use against an attacker is a “honey net,” which diverts attackers into a false network full of whatever the defender wants the attacker to see. If a cyberspace attacker is attempting to break into a highly classified system and the defenders know it, they can divert the attacker into a false network. Having blocked the attacker, there is nothing preventing him from trying again using a different access that the defender might miss. If the defender instead diverts the attacker but provides him with false information, it can be far more effective. For the defenders to be effective in their deception, they must understand the expectations of the attacker and provide an environment tailored to what the attacker expects to find.²⁶ Something similar to this may have happened in the early 1980s when a US spy provided information the Soviet Union was planning to secretly acquire gas pipeline technology. Instead of blocking the sale, the United States allegedly quietly altered the computer code that eventually led to “the most monumental non-nuclear explosion and fire ever seen from space.”²⁷ Defenders can do much the same thing in cyberspace. Once a defender captures an attacker in a honey net, the defender can keep the attacker busy with false information, examine attack patterns and techniques, embed beacons to phone home in the data that the attacker is taking, or carry out whatever other countermeasures are most useful. One of the most worthwhile techniques for a defender is instilling doubt in the mind of the attacker.

Introducing doubt into the mind of an attacker is one of the more useful things a resilient cyberspace defender can do with a honey net. A defender does not have to falsify everything; successfully falsifying one piece of information can make the attacker doubt everything else he or she got as well. One way of accomplishing this increased doubt is through a defender falsifying battle damage assessment (BDA). It is normally difficult for an attacker to understand how effective her or his attacks have been; a honey net can make it even worse. A defender can use a honey net to make it look like an attack has been successful, but then suddenly turn the system back on to ambush the attacker’s forces at the most opportune time.²⁸ Tricking the enemy this way once will also have the effect of making the enemy very reluctant to trust any future cyberspace BDA. If an adversary does not fall for a honey net,

a resilient cyberspace defender should have multiple copies of critical data available.

Backups enable a defender to rapidly reconstitute damaged systems and data and provide a way for defenders to minimize the effect of even a successful attack.²⁹ An attacker who breaks into a logistics system and erases all the data can cause significant problems for a defender. However, if the defender has a backup and can have the system restored and operating in a day, the defender can minimize the attack's long-term effects. One of the hopeful trends for cyberspace defenders is the decreasing cost of electronic data storage. This decreasing cost makes it far easier for defenders to keep multiple copies of the data needed to restore a system. Of course, defenders need to keep the copies in a manner that prevents an attacker from getting to the backups and the primary system at the same time.³⁰ Automatic backup systems may be convenient, but they are automatic and will copy a cyberspace weapon just as easily as valid data. Backups are part of resilience, but cyberspace defenders can also build hidden additional capability into their networks.

A war reserve mode (WARM) is a concept from electronic warfare that has great applicability in cyberspace combat. Electronic warfare equipment, such as radars or radios, is often built with additional functionality that is not used unless needed in a major conflict against a top-tier enemy. The reason for these hidden modes is that every technique has a countermeasure, and if all of a combatant's techniques are used routinely, the enemy will find out what they are and develop countermeasures. If the best techniques are hidden and not used until combat starts, it can give one side a decisive advantage.

Applied to cyberspace, WARM would suggest that defenders have preplanned ways to significantly alter not only their defenses but also their networks in ways that make an attacker's careful reconnaissance obsolete. As Gregory Rattray, a former US Air Force commander for information warfare and a cyber expert at the National Security Council, has stated, in cyberspace the equivalents of mountains and oceans can be moved with the "flick of a switch."³¹ Additionally, defenders do not have to accept the geography of their environment; they can actively change the terrain to make it harder to attack.³² Cyberspace attackers often have to spend significant time and effort mapping out exactly how a network is configured and what software it is running. If a defender was to change all the software on his routers to a previously unknown

version just as a conflict starts, it could disrupt many of the attacker's plans. The new software does not even have to be better than the old one or have less vulnerabilities. The new vulnerabilities will still have to be discovered, which can buy the defender significant time and breathing space. The same principle could be utilized for any software on the network and even has applicability to hardware.

A well-resourced defender could have significant spare hardware of different types on hand to enable a quick rebuild of the network. For example, if a defender has a diversified network, with three types of routers, if one of those routers is successfully attacked, the defender could replace the vulnerable one with one of the other two types from storage. Defenders could go so far as to build entire networks using different types of hardware that are then left in a standby mode and disconnected from other systems, which makes them very difficult to attack. If the primary system is successfully attacked, the defender can switch to the backup. Defenders cannot simply set up a backup system and assume it will work when needed; they have to extensively test and evaluate it. Otherwise, a backup system may be useless, as it provides a false sense of security but no capability when it is needed.

Of course, setting up an entire backup network is extremely expensive and will likely only be worthwhile on a small scale and when the information or network is so critical that it cannot be allowed to fail. Nuclear command and control is one obvious area where the requirement for surety is so high that a complete backup network is reasonable. These techniques of improved situational awareness, effective command and control, hack backs, honey nets, backups, WARM, and backup networks will help cyberspace defenders dynamically maneuver and bend in the right direction when the attacking wind comes.

Conclusion

There are many ways cyberspace defenders can bend under attack before springing upright like the bamboo in the Japanese proverb. The three different elements that apply to the resilience of the bamboo as well as resilience in cyberspace are flexibility, reducing attack surfaces, and reacting dynamically to attack.

Flexibility in cyberspace conflict will largely stem from creating flexible people, although good network design that allows flexible cyber-

space operators more options will also help. Flexibility in cyberspace systems will be expensive, and efficiency will be lowered due to the necessary excess capacity that must be built into a more flexible system. A flexible cyberspace system is also one that is heterogeneous and broken into defensible enclaves, not one large and easy to administer network running the same software on every device. Flexible cyberspace personnel are best grown through extensive education, training, and exercises, including red teaming, full-scale exercises with cyberspace play allowed to affect the physical domains, and accountability for users who prove unable to adopt good security practices. Many of these changes will be very hard to implement, as they will be extremely uncomfortable to bureaucratic organizations and will require significant cultural change.

Reducing attack surfaces is the second element to creating cyberspace resiliency. The first and extremely difficult step is to eliminate unnecessary capability across the network, both in software and in hardware. Users will be discomfited by having to use different tools than they are used to, but the payoff in security can be significant. Not every backup system should be eliminated. Wherever possible, a primary and backup for each key mission area should be available to allow cyberspace operators to rapidly shift from one to the other if vulnerabilities are discovered. Users' desire for continual rapid improvements in communication and capability will also have to be balanced against security requirements, as each new capability or communications pathway introduces a potential attack vector. Striking the correct balance between security and capability will be a difficult and continuing challenge, and the correct balance will change depending on the organization and environment within which it operates.

The final element of resiliency in cyberspace is the ability to react dynamically to attack. Cyberspace operators need to develop better situational awareness of their own networks and develop intelligence capabilities to understand what the enemy is planning. Attackers and defenders on the same side also need to lower the walls between them and share more information to enable cyberspace defenders to be better able to defend their networks, while protecting offensive capabilities. Effective cyberspace command and control that treats cyberspace operators as maneuvering forces to be commanded versus an IT management problem with an engineering solution is also important. Active defenses,

including honey nets, backups, WARM, and backup hardware contribute to cyberspace resilience.

For an organization to create enduring cyberspace resilience, some aspects of all three elements will be required, and building in cyberspace resilience will not be cheap. The additional costs incurred for redundancy and training alone will overwhelm any potential savings from streamlining and reducing excess capacity. However, if an organization is serious about protecting its ability to accomplish its mission in cyberspace, resilience under attack will be the key. If cyberspace operators and their defended networks and systems adopt the characteristics of the versatile bamboo, they too will be resilient enough to spring upright after the passage of the storm. **SSQ**

Notes

1. Risk Steering Committee, *DHS Risk Lexicon: 2010 Edition* (Washington, DC: Department of Homeland Security, September 2010), 26, <http://www.dhs.gov/xlibrary/assets/dhs-risk-lexicon-2010.pdf>.
2. Peter W. Singer and Allan Friedman, *Cybersecurity and Cyberwar: What Everyone Needs to Know* (New York: Oxford University Press, 2013), Kindle location 720. Singer and Friedman have recently suggested that the classic information security “CIA Triad” of confidentiality, integrity, and availability should be extended to include resilience.
3. Martin C. Libicki, *Cyberdeterrence and Cyberwar* (Santa Monica, CA: RAND Corporation, 2009), 176.
4. William D. Bryant, “Cyberspace Superiority: Dominating the Digital Frontier” (PhD diss., School of Advanced Air and Space Studies, 2014), 205.
5. Bryant, “Cyberspace Superiority,” 206. I was only able to examine and code eight cases, due to limitations in the available unclassified data. The four successful cases in descending order of the level of success were Russia v. Georgia in 2009, the Nonghyup bank attack in 2011, Stuxnet, and the South Korean distributed denial-of-service (DDoS) attack of 2011. The four unsuccessful cases in decreasing order of the level of success were Russia v. Estonia in 2007, Aramco in 2012, the 2013 bank and media company attacks against South Korea, and North Korea’s 2009 DDoS attack against the United States and South Korea. The evidence and measurement methodology supporting this coding can also be found throughout my cited dissertation.
6. Joint Chiefs of Staff, Joint Publication 3-13, *Information Operations*, 27 November 2012, II-9.
7. Edward N. Luttwak, *Strategy: The Logic of War and Peace* (Cambridge, MA: Belknap Press, 2003), 39–40.
8. Aristotle, “Nicomachean Ethics,” in *The Book of Virtues*, ed. William J. Bennett (New York: Simon & Schuster, 1993), 102.
9. Paul Rosenzweig, *Cyber Warfare: How Conflicts in Cyberspace Are Challenging America and Changing the World* (Santa Barbara, CA: Praeger, 2013), Kindle location 3727.
10. Antoine Bousquet, *The Scientific Way of Warfare* (New York: Columbia University Press, 2009), 222.

11. Gregory J. Rattray, *Strategic Warfare in Cyberspace* (Cambridge, MA: MIT Press, 2001), Kindle location 220.
12. David J. Lonsdale, *The Nature of War in the Information Age* (London: Frank Cass, 2004), 154.
13. Gen James E. Cartwright, US Marine Corps, comments, Air Force Association Air Warfare Symposium, 8 February 2007, reported in Franklin D. Kramer, "Cyberpower and National Security: Policy Recommendations for a Strategic Framework," in *Cyberpower and National Security*, edited by Franklin D. Kramer, Stuart H. Starr, and Larry K. Wentz (Washington, DC: Potomac Books, 2009), 14.
14. Rosenzweig, *Cyber Warfare*, Kindle location 815, chap. 3.
15. Rattray, *Strategic Warfare in Cyberspace*, Kindle location 219.
16. Martin C. Libicki, "Cyberspace Is Not a Warfighting Domain." *IIS: A Journal of Law and Policy for the Information Society* 8, no. 2 (Fall 2012), 330.
17. Gilmary Michael Hostage III and Larry R. Broadwell Jr., "Resilient Command and Control: The Need for Distributed Control," *Joint Forces Quarterly* 75, no. 3 (2014): 38–43.
18. William A. Owens, Kenneth W. Dam, and Herbert S. Lin, eds. *Technology, Policy, Law, and Ethics Regarding U.S. Acquisition and Use of Cyberattack Capabilities* (Washington, DC: National Academies Press, 2009), 84.
19. Jeffrey Carr, *Inside Cyber Warfare: Mapping the Cyber Underworld* (Beijing: O'Reilly Media, 2011), Kindle location 3674.
20. Libicki, *Conquest in Cyberspace*, 74.
21. Carl von Clausewitz, *On War*, edited and translated by Michael Howard and Peter Paret (Princeton, NJ: Princeton University Press, 1976), 75.
22. *Ibid.*, 357.
23. Owens, Dam, and Lin, *Technology, Policy, Law, and Ethics*, 13.
24. Rosenzweig, *Cyber Warfare*, Kindle location 2024, chap. 7.
25. Kelly Jackson Higgins, "Free 'Active Defense' Tools Emerge," *Security Dark Reading*, 11 July 2013, <http://www.darkreading.com/intrusion-prevention/free-active-defense-tools-emerge/240158160>.
26. Colin S. Gray, *Modern Strategy* (New York: Oxford University Press, 1999), 35.
27. Thomas C. Reed, *At the Abyss: An Insider's History of the Cold War* (New York: Ballantine Books, 2004), 269. Other authors, such as Thomas Rid, have questioned whether the explosion actually occurred. Reed, given his various government positions, including Secretary of the Air Force, was in a good position to know of the events if they occurred as alleged. See Thomas Rid, *Cyber War Will Not Take Place* (New York: Oxford University Press, 2013), Kindle location 275.
28. Owens, Dam, and Lin, *Technology, Policy, Law, and Ethics*, 125.
29. Libicki links the replicability of cyberspace with its reparability in *Conquest in Cyberspace*, 5.
30. Singer and Friedman, *Cybersecurity and Cyberwar*, Kindle location 3177.
31. Gregory J. Rattray, "An Environmental Approach to Understanding Cyberpower," in *Cyberpower and National Security*, edited by Franklin D. Kramer, Stuart H. Starr, and Larry K. Wentz (Washington, DC: Potomac Books, 2009), 256.
32. Libicki, "Cyberspace Is Not a Warfighting Domain," 324.

Disclaimer

The views and opinions expressed or implied in SSQ are those of the authors and are not officially sanctioned by any agency or department of the US government. We encourage you to send comments to: strategicstudiesquarterly@us.af.mil.