CHAPTER 7

Strategic Defensive Security

In the previous chapter, we discussed a scenario in which a CISO was given a budget of $5 million to develop the cybersecurity program of a sneaker company. Our CISO has been given a very large task, but before they begin, they need to consider the objective of their security program. The objective may seem obvious, in fact, it is so obvious that it is rarely even considered worth mentioning or debating – don't get hacked! Cybersecurity means preventing the bad guys from getting in, so this program should do just that, right**?**

This is what I would consider the classical approach to a cybersecurity program. This principle has been the core of the cybersecurity conversation for decades, with CISOs and their security personnel attempting to prevent any sort of compromise across their organization. Resources are spread wide, security solutions deployed equally on every host, and analysts watch for any sign of compromise across thousands of network points, turning cybersecurity into a vast game of whack-a-mole played at the speed of processors.

But before we go down the road of the classical approach to cybersecurity, perhaps we should take a moment to learn the lesson that cybersecurity has been trying to teach us for decades – it doesn't matter what technology you use, how many highly skilled security personnel you employ, or how locked down your policies are, you will get hacked.

J. G. Oakley et al., *Theoretical Cybersecurity*, https://doi.org/10.1007/978-1-4842-8300-4_7

An adversary with time, resources, and motivation will find a way into any network regardless of the size of the cybersecurity budget. This principle is entirely counter to the classic approach. Strategies that seem obvious under the lens of the classic approach now seem outdated and clearly flawed. Attempting to secure every element of your IT infrastructure to the same level doesn't make as much sense when you no longer believe it is possible to not get hacked. However, I would argue that this principle is more mature and results in a significantly more effective security program.

This principle is the foundation of the strategic defensive security approach and in this chapter, we will examine several different aspects of a well-built security program and contrast the classic approach with the strategic approach as our CISO is determining their approach to the sneaker company's cybersecurity.

Architecture

Just like any New York skyscraper, a strong cybersecurity program begins with the architecture. The architects are responsible for designing how the system will be laid out, the broad strokes of the implementation on a technical level, and the phases of capability construction. Architecture is where the priorities of the CISO become clear both in terms of technology and budget share as each are divvied up across the organization's IT landscape.

The Classic Approach

A CISO using a classic approach to their architecture is going to have one primary, often unspoken, goal – don't get hacked. Once this golden principle is accepted, the next step for our sneaker company CISO is to set the priorities for the various aspects of the IT footprint in terms of budget shares. While there are any number of ways to break down these

decisions, the simplest approach is to view the business network in terms of internal and external. The internal network is considered to be the more sensitive side of the network. It is where internal processes are executed and generally is not accessible via the Internet without some form of authentication that hopefully prevents unauthorized access. The external network faces the Internet and provides the public with access to whatever applications, store fronts, etc. that are necessary to support the public offering of the company.

If our CISO's goal is to prevent any attack from becoming a full-blown compromise, then securing externally facing assets seems like the best place to start. Our CISO therefore prioritizes any assets that can be reached by a threat originating somewhere out there in the Internet. Our CISO reasons that these assets will be expected to weather the majority of attacks, they need to be the most secure resources within the organization. They are viewed as the wall the attacker must climb to get a glimpse of the more sensitive internal network.

Internal assets are not as high of a priority since the CISO, and his team of architects do not expect them to face as many threats. That is not to say that they are entirely neglected. The internal network will be secured as far as its lesser-prioritized budget will allow. For example, our CISO could obligate $2 million of the available $5 million for internal security, granting the larger share to the external side.

The classic approach results in a cybersecurity construct that very much resembles a medieval castle. The walls are large and thick. The defenders of the castle are perched on top of the walls waiting to shoot down any approaching attackers. Any attacker approaching the castle along the expected routes will find themselves intimidated by the defenses they are attempting to overcome. However, if one was able to view the castle from above, the security begins to show some weaknesses. Yes, the walls are big and strong but anyone who breaches them will find that there are few if any defenders within.

This classic approach has its flaws and from this overhead view, you may be already spotting them. If you've worked within the cybersecurity field for a few years, you may find this example contrived, overly simplistic, and yet…familiar. Don't blame our poor CISO. We can follow the logic; the path that leads from the core principle of "don't get hacked" to the emphasis on external asset protection to the consequential de-prioritization of the internal network and subsequent trust.

This approach is the reality for many companies of every size. During my experience on the offensive security side of the spectrum, my teams have referred to these networks as the "gooey center." All an attacker has to do is crack or get around that strong external shell and the internal network becomes a free-for-all. The security of the external assets will raise the bar for the skill required to compromise the network which, of course, will reduce the total number of compromises. However, when a compromise does occur, it has a much higher likelihood of being catastrophic.

The Strategic Approach

Our sneaker company CISO has been around the block a few times. He's seen the kind of security programs that are built on the golden principle and he isn't a fan. He opts to do away with the golden principle and instead starts with a different principle – "we will be hacked." Like we saw in the classical approach, the logic of prioritization will most often follow a natural path that originates and is based upon the guiding principle. But this time the guiding principle is different, and it will result in a different logical flow.

Our CISO calls a meeting with his team of security architects. He writes on the board, "We Will Be Hacked" before taking his seat. "I want to build our security program with the expectation that we will be compromised. Our goal is not to prevent every compromise, our goal is to develop architecture that will withstand a compromise without allowing a

disruption to critical business operations. We will be compromise-resilient and we will render compromises inconsequential even though they will occur."

The ideas seem radical, but the team gets to work developing and then attempting to answer the questions that logically flow from this principle. What does this kind of network look like? How can network and security architecture be used to ensure continued business operations during a compromise?

Through the next few days and perhaps weeks, our CISO works with his architects. First, they determine that in order to ensure continued business operations, they must protect assets that are critical to those operations. That might include the primary public facing ecommerce site where their customers go to buy their sneakers. It might include the databases that hold sensitive client payment information for recurring purchases (only for the most dedicated sneaker-head). It might include the backend processing systems that allow credit card processing to occur. The team realizes that if they are to use the guiding principle of expecting a compromise, they must begin by prioritizing assets by their ability to affect critical business goals.

A layout begins to take form. A layout that looks like an unfinished connect-the-dots puzzle, wherein the dots represent network nodes considered critical to business goals. During this process, one architect examines the layout and speaks up. "Wait a minute. Where is the internal and external boundary line? How will we know which nodes we need to protect more if we don't know whether they are publicly or privately available?" The other architects examine the layout and consider the question. They realize that while in practice, defending public and private nodes will be different since they will not experience the same type of attacks, from a prioritization point of view, there is no difference. Technology and budget-share prioritization will be given to these critical nodes regardless of which side of the network they fall on.

The team quickly realizes that simply identifying critical nodes is not enough. The dots are not isolated. There are network paths and other hosts that connect them and, if compromised, these connections could also threaten business operations. However, these connections are not quite as critically important as the critical assets themselves. The team determines that the hosts directly connected to critical assets should be labeled as High in the criticality scale.

The team continues this process. The logical layout of network security is constructed like ripples in water. Critical assets are at the center with assets of lower criticality levels encircling them and expanding the further away from the critical asset they are.

A full week into the development of the architecture, a senior architect notices something and raises her hand. "There's too many connections. Half our network is rated at High criticality. We can't focus budget share and effort on half the network!" The other architects examine the designs and are forced to agree. Proximity to critical assets is everywhere. Afterall, as critical assets, they hold data that is going to be used by much of the network. The architect has an idea, "We need to reduce the connections; isolate the critical assets as much as possible. We'll need to construct connections to these critical assets with very tightly defined access and focus much of our defensive capabilities on those connections. That way, we will prioritize the critical asset as Critical, the connections as High, and the assets using the connections as Medium or perhaps even Low, depending on their ability to access levels of sensitive data."

After weeks of hard work, the team emerges exhausted and holding a plan that does not prioritize one large section of the network over the other. In fact, internal vs. external conversations were avoided entirely. The final plan contains a chain of assets that make up the most critical infrastructure for the continued operation and stability of the sneaker company. These assets will be prioritized over all others so that in case a compromise does occur, the core business assets will weather the storm.

These critical assets include the primary ecommerce website of the sneaker company, supporting databases with sensitive client data within the internal network, credit card processing servers, sneaker design and other intellectual property storage, and more.

The critical assets are prioritized regardless of their place in the network since a compromise of any asset in this chain would have a critical effect on the company objectives. The critical chain receives the necessary portion of the budget to ensure the architects are able to lock it down at a level commensurate with its priority. The next priority of assets is those that can affect business operations at a High rather than Critical level or those that have close access to critical assets and could be used to break into the critical chain. Outside of High are the Medium level assets and so on, with each level of assets receiving less budgeting and manpower prioritization.

The architects set about securing critical assets first, keeping in mind that assets that are not members of this chain have an expected higher chance of compromise. However, a compromise of these assets would have a less significant impact on core business objectives than the assets deemed critical.

Monitor and Detect

Both versions of our sneaker company CISO have completed their architectural plans. The classic version has a standard network focused on external security and a "don't get hacked" mentality, while the strategic version has a network focused on the security of what matters to company operations and a "we will get hacked" mentality.

Both versions now turn their attention to the next part of their security program – security monitoring and detection.

The Classic Approach

Our classic approach CISO begins constructing his security monitoring and threat detection program based upon the golden principle. Again, this principle is never really spoken. All members of the security team implicitly know that their mission is to simply not get hacked. Monitoring and detection capabilities will be focused onto that idea from the ground up.

As our CISO sets to work, they quickly discover a new principle as a logical result of the golden principle applied to monitoring and detection: Visibility is king! All IT assets must be monitored regardless of their location. Internal, external, cloud-based, or the break room smart fridge, everything must be monitored. Logs must be aggregated from every potential source so that detections that sweep across the entire network can be written. The core idea being that if you can't see it, you can't tell if it's under attack or gasp compromised! A network must eliminate blind spots to eliminate the threats lurking therein.

Our CISO performs research on the topic and finds himself in good company. Cybersecurity leaders reinforce his idea of the importance of visibility across the network. White papers have been written pushing the theory and building upon it.

With confidence in his approach, the CISO builds a list of the latest security features and solutions that he believes will best defend his network. He researches additional add-ons for solutions that will help automate the response to detected malicious activity. Finally, he begins the real-world implementation of the program by reaching out to the security solution vendors.

It is at this point that we spot the flaw in the classic approach. It is an item that all security programs grapple with, nearly all security professionals complain about, and the core reason that we can't make things as secure as we want to – Budget.

The CISO begins to review the cost of his approach and finds that the dominating pricing model for all aspects of his security monitoring and detection program is that they scale with size. The larger the network, the more they cost. Logging and aggregation solutions cost more as they process more data. Endpoint protection and monitoring cost more as more endpoints are protected. The add-ons to the given solutions that the CISO felt would greatly increase the network's resilience to attack add additional expenses, and once again, these expenses scale.

The CISO finds that he is forced to reevaluate his list. He trims down some of the more expensive plugins and selects less costly log aggregation solutions. He simply cannot afford to deploy the level of security solutions he would like in his network due to its size and must make compromises somewhere in order to achieve visibility across his network.

Our CISO ends the construction of his monitoring and detection program feeling rather depressed about the fact that budgets limited his ability to implement his golden principle. In the back of his mind, he realizes that by requiring every aspect of his network to have the same level of protection, while also being unable to pay for the level of protection he would have liked, his entire network is now less secure than he would have liked.

The Strategic Approach

Our classic CISO had lofty goals rooted in the best ideas for secure security program. It could be summarized as the philosophy of log, monitor, and detect everything. Unfortunately, it could not stand the reality of budget limitations and the very real-world effect they have on such goals.

The CISO using the strategic approach already has a vehicle to enable budget limitation considerations. That vehicle is the prioritization of assets with the network that was established during the creation of network security architecture. The philosophy of the strategic approach to security monitoring and detection could be summarized as log, monitor, and detect what matters.

The CISO goes through the same basic process that the classic approach CISO undertook. He begins by developing a list of security solutions that will best secure his network. It's an expensive list of some of the leading products complete with cutting-edge upgrades and add-ons. It's a list that would eat through his budget in a heartbeat if applied to the entire network. The CISO writes the world "Critical" across the top of that list and then sets it aside. He begins to write a second list. This list contains technologies that are slightly less featureful but still powerful. As you might guess, he writes the word "High" across the top of that list. A list is developed for every criticality level with decreasing features and a corresponding decrease in cost.

With these designs, the engineers set to work implementing the various levels of monitoring and detection products across the levels of criticality. Since the CISO has taken budget limitations into consideration at every level, he is not surprised by the final product. When the dust clears, the CISO finds that he has incredible capabilities allowing his team to protect the most critical assets.

Our CISO is well aware that this approach means that aspects of his network will be more "in the dark" than others. There will be places where a compromise could occur and not be immediately detected. However, his focus is on maintaining the strategic goals of the sneaker company. He is enabling the continued operations of critical business goals even in the face of a compromise. He is focused on ensuring that compromises, when they do occur, are not able to breach the upper echelons of network criticality.

I recognize that the ideas in this approach are controversial. Egalitarian visibility and monitoring are the core of modern security architecture. In theory, I entirely agree with this concept. If budgets were not a concern, then more data is always a good thing. Unfortunately, budgets are a concern, in fact, they are the primary limiting factor. Networks are not equal. Certain assets have a far greater ability to affect the overall mission of the organization than others. Ignoring this fact results in a network

made less secure due to the ratio of its size against its budget. By securing everything to the same level, we lower the level to which everything is secure.

Investigate

Cyber investigation, more commonly called "hunting," is the process of proactively examining data within a given network for evidence of a compromise. It is an evolved approach that recognizes the shortcomings of standard reactive SOCs who are only aware of a compromise if an alert has been written for the actions that the attacker has taken. Threats evolve and organizations have to find ways of responding to new attacks and new types of compromise. Cyber investigation is a step in that direction.

Both versions of our CISO, the classic and the strategic, see value in the proactive approach of investigation and set out to develop the goals and guidelines of an investigation program.

The Classic Approach

Our golden principle-minded CISO isn't feeling the best after budget limitations derailed some of his goals for his monitoring program, but he discovers he still has some money left over for an investigation program. He is excited. He has heard so much about the developments in cyber hunting and threat intelligence and he can't wait to implement some of these new ideas in his environment.

He begins to build his team based on the same approach that many of his peers are using. Like much of cybersecurity, cyber investigation has largely been developed into a generalized "one size fits all" approach. Investigators (or "hunters") examine threat intelligence streams to learn trends in current threats and develop automated means of examining logs for those behaviors. Well-known behaviors of documented threats

are recorded in standardized frameworks like the MITRE ATT&CK framework and investigators rely on them to hunt for behaviors of known threats. Threat intel streams and frameworks are often broken into broad categories of sectors such as federal, local, commercial, non-profit, financial, etc.

Our classic CISO instructs his team of investigators to ensure that all known actor techniques are easily detected by the monitoring capabilities. Threat streams are purchased and watched so that investigators know what new threats are occurring in the market in general and can ensure that the known behaviors for those sectors are detected.

The generalized approach may seem great for those that subscribe to the "don't get hacked" philosophy. After all, what harm could come from protecting yourself from the behaviors of all types of threats, even if they don't always apply to your industry? Organizations using this approach purchase membership to several threat intel streams and ensure that they are capable of detecting behaviors documented in hunting frameworks, assured that they are at least as aware of threats as other similar organizations.

The issue in this approach comes from the generality that exists at its core. Smaller organizations can use a generalized approach because they are at much less risk of being directly targeted. It is less likely that a somewhat sophisticated actor is developing target attacks against them. However, as an organization grows, the generalized approach to investigate becomes an ever-expanding blind spot. This blind spot exists in two ways. First, the organization is not aware and not prioritizing the specific attacks that are targeting them due to their "don't get hacked" approach, and second, the organization is not aware of the value a threat places on compromising their networks. This second concept takes a bit of unwrapping, so we'll examine it in detail in the next section.

The Strategic Approach

Like the classic CISO, the strategic CISO is also excited about cyber investigations and threat hunting. The difference between the two again comes down to the difference in their founding principles. Where the classic CISO used a general approach to his investigation program, the strategic CISO is interested in focusing his investigation on the specific threats that his organization faces.

Before he gets too deep into investigations, our strategic CISO wants to first determine the value of the assets he is trying to protect. That might seem pretty straightforward. If the CISO wants to protect the intellectual property for a sneaker that is currently being designed, for example, the value could be easily derived by determining how much revenue would be lost if the design were to be tampered with or leaked. The CISO could logically conclude that he should not spend a portion of the security budget protecting the shoe design that is greater than the amount of revenue that would be lost if the design were compromised.

There is a significant assumption at the core of this approach to valuation. This assumption represents a flaw in both the valuation process and the classic approach to threat investigation. The assumption is that the threat will value an organization's assets in the same way the organization does. From a security perspective, an asset's value is not solely determined by what an organization stands to gain from the asset. It is also determined by the attacker.

Let's take a break from the world of sneakers to explore this a bit further through a real-world scenario that I experienced during my career. I was performing security assessments and penetration tests for a Fortune 100 company. This company developed many widely used applications and the budget that was set aside to defend these applications was built on a given application's worth to the company. However, the company begins to encounter incredibly sophisticated attacks beyond the capabilities of the budget they had provided. They determined that foreign governments

had a deep interest in attacking the company. At first, the company was confused. They were just a commercial company developing products like any other software company. Why would they be the target of a sophisticated nation state adversary? The answer was that the products the company developed were more valuable to the nation state adversary than they were to the company. The nation state wanted to obtain source code for those products to identify vulnerabilities since the US government widely used the products. The company was not aware of the kind of threats they faced. They used the same threat intel streams as others in the same market. They used the same generalized approach instead of working to understand the specific threats that they faced. They used their own valuation of their products and allowed that to drive their defense instead of determining how much their product was worth to an adversary.

General knowledge about the cyber threat landscape is very helpful, but it isn't the whole picture. An organization needs to identify the specific threats it is facing and understand the level of resources that threat is willing to spend compromising the organization. As an organization grows, general threat and adversary technique knowledge becomes less and less useful since threats have an increasingly varied set of reasons for attacking it.

Getting back to our CISO, he believes that his organization is large enough that an examination of more specific threats is in order. His team of investigators returns after sometime to inform him that the sneaker designs are actually quite novel in the market place and have a chance to revolutionize a section of the market. Competitors are very interested in the intellectual property and may even be attempting various forms of corporate espionage to obtain the designs. The CISO realizes that protecting the shoe designs will require more of a budget than he had originally considered through his revenue-based valuation.

Frameworks

If you've worked in cybersecurity or managed virtually any application that was used by the federal government, the military, processed credit cards, or something else considered sensitive, you've undoubtably encountered certification frameworks. These frameworks are created by large government, military, or commercial entities to ensure that a product, application, network, etc. meets a minimum security standard. If you've worked with credit card processing applications, then you've had to deal with PCI certification. If you've attempted to sell cloud-based applications to the US federal government, you've encountered FedRAMP. The FedRAMP process is illustrated in Figure 7-1.

Figure 7-1. *FedRAMP Authorization Process*

Certification frameworks are an important part of ensuring a standard level of security before an application is trusted with some level of sensitive information. Unfortunately, these frameworks do little to ensure the resilience of the applications and networks they are applied to.

We will take the example of the Federal Risk and Authorization Management Program, more commonly called FedRAMP. FedRAMP is a certification framework specifically applied to cloud-based offerings and is used to ensure a level of security among these offerings as they are sold to the United States federal government. FedRAMP is one of the most modern certification frameworks. In my opinion, it does a pretty good job. Its requirements are more in-depth than most frameworks and more aware of the intricacies of the various cybersecurity disciplines they span.

For all the forward thinking of the FedRAMP requirements, it is still based on that golden principle of the classic approach – don't get hacked. It has no requirements regarding the resiliency of the application nor any requirements that dictate how data exposure could be minimized to help render compromises inconsequential. Like most of the cybersecurity industry, the FedRAMP framework does not consider that a portion of the application that attains its certification will inevitably be compromised at some point.

When we examine FedRAMP in the context of expecting a compromise regardless of how good the standards are, the blind-spot of resiliency becomes obvious. If security is only half the battle, FedRAMP is only half a framework.

We can apply this context to practically every other certification framework that is similar to FedRAMP. The Risk Management Framework (RMF) was originally created in 2004 by NIST and then updated in 2018 through NIST SP 800-37. This framework is used by every agency of the US federal government and the DoD. It defines a high-level seven-step process for securing systems through an Authorization to Operate (ATO) and ongoing risk management, often referred to as continuous monitoring. The intent of RMF is to be technology agnostic so that it can be used to apply security and risk management at every level.

You can problem-guess what's coming. RMF is built on the golden principle and lives in a world where compromises do not happen if security controls are tight. Its seven steps can be summarized as

1. Prepare
2. Categorize information systems
3. Select security controls
4. Implement security controls
5. Assess security controls

6. Authorize information system
7. Monitor security controls

Figure 7-2 illustrates these steps in a commonly shown depiction of the RMF process presented by NIST and other organizations that implement it.

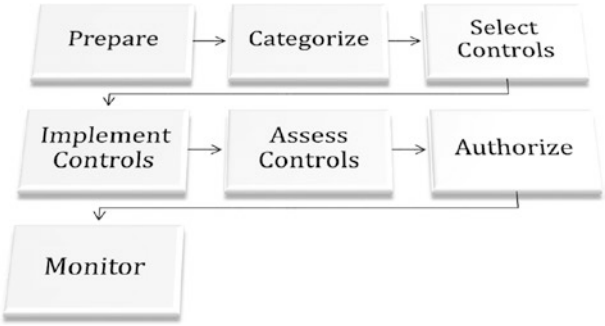

Figure 7-2. *RMF Steps*

What's missing here? Perhaps an eighth step called "Simulate security control failure," and a ninth step called "Minimize compromise impact"? Again, if we come from the understanding that a portion of the systems secured through RMF will still be compromised, then suddenly the steps as they are established by NIST seem like an unfinished sentence.

I won't bore you by examining every certification framework. Rest assured that CNMC, PCI, and others suffer from the same lack of resiliency consideration. They are frameworks that are missing half the battle.

So what would a framework that considers both security and resiliency look like? We are already beginning to see some movement in that direction within the industry. These movements have been spurred on by the rise in ransomware attacks. The term "ransomware attack" is interesting. Ransomware is not an attack. Ransomware is a payload executed after a successful attack. This differentiation is important because it shows that ransomware is only highlighting the compromises that were already occurring. Our applications were already compromised, our frameworks were already failing. Ransomware just turned up the stakes.

Let's return to the FedRAMP discussion. How could FedRAMP be changed to consider the importance of resiliency and compromise survivability? Again, we begin with the idea that a compromise will occur. If we accept that fact, then the next step is to gain an understanding of the consequences of compromise at different levels. In addition to its enforcement of security standards, FedRAMP-certified applications should provide an impact analysis for production server compromise, database compromise, cloud account compromise, etc. The given applications should attempt to be as secure as possible but also grant their federal customers an understanding of the exposure at these various levels of compromise.

In the current FedRAMP framework, vulnerabilities that are identified through the assessment of security controls are rated High, Medium, or Low severity and given time windows for remediation. If the vulnerability cannot be remediated due to the functionality's importance to the overall product, the certifying federal agency must either accept the risk or reject the product. This concept could be applied to the resiliency side of the framework. The application seeking certification would provide data on the exposure that results from various compromise scenarios. If the exposure is unacceptable, the company selling the application must work to minimize the exposure and increase the resiliency of its product.

FedRAMP is a convenient example, but the focus on resiliency as an equally important objective as security can be woven into any cybersecurity framework.

Auditing

A framework is only as good as the standards of its audits. RMF, for example, was intended to be a flexible framework applied at any level without consideration for the specific technology it was applied against. However, in practice, RMF can become little more than a checklist. Its

deeper implications and intentions can be lost when auditing practices are not firm, documented, and enforced.

Auditing within a framework that implements resiliency would require the simulation of compromise scenarios and then examination of what data or impact those compromises can affect. FedRAMP already requires an in-depth review of all security controls and their implementation and even establishes thorough requirements for a penetration test. If FedRAMP were to be expanded to consider resiliency then a resiliency test, similar to a penetration test, would be included. Different levels of compromise would be created within the application seeking certification and the auditors would examine what data is exposed. Risk ratings would be applied to the exposure and the company developing the application would be required to implement better resiliency policies to reduce the risk ratings and obtain certification.

Theoretical Case Studies

So far, we've explored the concepts of strategic defensive security within the context of commercial companies almost exclusively. In the next section, I'd like to demonstrate how the same concepts can be applied to other sectors through the use of hypothetical case studies.

The Architecture of Accountable Sectors

The Springfield Children's hospital has discovered that it is the victim of a ransomware attack. Five doctors are unable to access their data and treat their patients. The ransomware demands a payment in crypto currency. After a brief meeting, the hospital directors pay the ransom and re-gain control of the computers.

This experience has shaken the trust that the directors have in their IT infrastructure. They call in their director of network security and ask how they can be sure that these kinds of attacks will never be successful again. The director simply states that they cannot be sure. In fact, similar attacks will most likely be successful in the future. The hospital directors task the director of network security with re-architecting the security program and, if necessary, the network itself to account for these kinds of attacks and to protect that which is most important to the hospital. The director of network security gets to work.

The director recently read a book on Strategic Defensive Security and decides to use that approach. The first step he takes is to define the mission objectives of the security of the hospital. From his research, he defines three such objectives and ranks them in order:

1. Protect patient lives
2. Protect patient health data
3. Ensure continued hospital operations

The next step is to identify the network nodes that have the ability to affect these objectives and assign the nodes a criticality level. The director examines network diagrams and identifies the systems directly responsible for the control of life support systems and marks them with a critical severity. Next, he notes any node which has the ability to affect life support nodes and marks them at a high severity. He continues through several rounds of increasing distance from life support nodes and corresponding decreases in criticality rating.

With network nodes associated with the first objective prioritized, the director moves on to nodes that hold sensitive patient health data. Through the same process, all nodes that have some proximity or ability to affect patient health data are prioritized. And finally, the process is repeated for the third objective.

The director pauses a moment and realizes that there are hundreds of network nodes capable of affecting these three objectives in some way or that are close enough within the network to affect nodes that could affect the objectives. He sets to work re-architecting the network with the goal of reducing the number of nodes with higher criticality rankings. He uses subnets, firewalls, and more to lock the most critical nodes off from the rest of the network except for defined access points. For networks relating to life support, the director splits them to their own network entirely with no connection to the general hospital network or the Internet.

With the number of nodes that need higher levels of protection reduced to the bare minimum, the director commissions his threat intelligence experts to create a profile of the kinds of threats the network will face. Their goal is to determine the level of value that threats place on hospital network. With this profile, the director will be aware of what parts of his network are valued higher by threat actors than the value the hospital itself might assign.

The security director examines his budget and selects cybersecurity products across a range of categories. Log aggregation, endpoint protection, etc. He implements the products with the greatest feature set on the most critical components and directs his Security Operations Center to prioritize events on those nodes above all others. Nodes with lower criticality ratings are assigned products with reduced cost as well as feature sets.

With the architecture and monitoring aspects of the network established, the director focuses on increasing the resiliency of the network to a compromise.

Military Resiliency

We've talked a little about resiliency and how necessary it is within the commercial sector. Within military sectors, it's a core requirement. Of course, the military isn't new to the concept of resiliency. Wars are messy

and combat in any form is disruptive and unpredictable by nature. So how does the military foster resiliency before engaging with an enemy? Training and experience are important factors but perhaps the most well-known test of resiliency outside of warfare is war games.

War games are a simulation of a battle. Both sides are staffed by military members attempting to outmaneuver the other. Creativity is encouraged and unexpected scenarios are guaranteed to occur. These games help commanders understand how to react when aspects of the infrastructure they rely upon are less than ideal or outright fail. But war games are used at every level. For example, the last step in the US Navy Boot Camp is a simulation called Battle Stations. During the 12-hour exercise, sailor candidates use their training to perform the mundane maintenance of a ship, while also responding to a number of catastrophic scenarios.

This style of building and evaluating resiliency is very similar to some of the ideas used within cybersecurity. Earlier we mentioned the Chaos Monkey project by Netflix which enforces resiliency by randomly shutting down servers within the Netflix production environment. Both of these approaches ensure resilience by creating unstable environments.

So how can we apply these concepts to military IT systems to foster and enforce resiliency at the level required by military objectives? If you have some experience within the cybersecurity community, you may be thinking that I'm about to discuss the common simulations that exist within the cybersecurity community today. These simulations are almost always Capture the Flag competitions, with defenders set on one side and attackers on the other. One well-known example of these CTF challenges is the National Collegiate Cyber Defense Competition wherein colleges compete to defend their simulated networks against trained penetration testers. These types of challenges are fun and have their place in cybersecurity education, but they are far from realistic. Cyber dogfighting across networks in real time is not a reality. CTFs should not be considered a viable means of learning resiliency.

Let's create a hypothetical case study. The military of Australia has realized that their IT networks were built with the outdated idea of "don't get hacked." Networks built with this principal are not resilient. They are built to prevent a compromise with the naive expectation that such a thing is possible.

Upon review, the Australian military leaders realize that they face advanced threats from around the world and compromise is inevitable. In fact, their systems are most likely already compromised to some level by the most advanced toolkits in the world. In addition, they are concerned that ransomware attacks and other debilitating threats could decrease the nation's preparedness for responding to a military threat.

The Australians decide to re-architect their cybersecurity program and part of that process is taking steps to create resilience within their IT networks so that a compromise is not able to significantly affect the military readiness. To accomplish this, they create cyber war game scenarios. The scenarios for these games include ransomware randomly deployed to a user's box, a domain controller compromised, an entire base losing Internet access, and more. These scenarios are executed as tabletop exercises at first, then they are conducted within test environments that mimic real military networks. But the Australian military leaders know that resilience is not achieved until the actual networks that are relied upon are put under the stress of cyberattack.

The Australian cyber command begins conducting simulated cyberattacks within the networks of various bases. The base commanders are given a warning that some level of attack will occur within a given window and a variety of compromise scenarios will be executed. In response to these cyber war games, new creative approaches are developed to maintain military IT objectives even during a significant compromise. Secondary fail-over networks are developed, sensitive data is available in fewer areas, and workstations are virtualized so that a response to a compromise can occur more quickly.