WHITE PAPER

Cyber Resilience – Healthcare's Most Important Asset

What Is Cyber Resilience?

It's no longer a case of if your hospital, clinic, insurance company, or other healthcare-related entity will be attacked — it's a question as to when. As a reference point, a Black Book Market Research report[1] showed that over 93% of healthcare organizations experienced one data breach between 2016 and late 2019 and 57% of organizations had over five breaches during that timeframe.

Organizations should be especially concerned about security breaches and ransomware in the current geopolitical climate. Attacks could come at any time. However, there is something you can do to mitigate breach risk. By adopting a resilient security architecture approach, the "time to observance" and "time to remediation" can be reduced. This is where cyber resilience comes into play. Cyber resilience is a strategy designed to help you reduce the cost and risk associated with a data breach.

A resilient approach allows you to:

- Strengthen your capabilities to defend against attacks
- Maximize your ability to rebound from an attack
- Minimize the severity and cost of security breaches

So, what do we really mean by resilience? Traditional resilience refers to the ability of an entity to return to its original form after being bent, stretched, or compressed. From our perspective, we are specifically talking about the ability of an IT network to recover to normal, steady state operations after a security attack or breach has occurred.

Network security resilience then is the set of activities that can be conducted to help the network after an attack happens. While most security architecture frameworks focus on preventing a breach, this security strategy is about "after breach" activities.

Cyber resilience deployment

A network security resilience concept has the potential to reduce the cost of a typical data breach. Faster detection and mitigation are key to controlling data exfiltration. Jon Oltsik, Enterprise Strategy Group.

[1] 2020 Cybersecurity Study, Black Book Market Research. 2020. https://blackbookmarketresearch.newswire.com/ news/healthcare-data-breaches-costs-industry-4-billion-by-years-end-2020-21027640

The reason for focusing on this strategy is simple — you want to reduce costs. These activities, if done right, will help you reduce the financial and emotional costs of a breach.

So, to be clear, we're not suggesting that you stop trying to prevent a breach. You always want to do that. What we are saying is that you want to add another set of capabilities to lower your corporate risk and the cost of a breach. At some point in time, you are going to be hacked. The question is, how painful and expensive do you want that breach to be?

The Importance of Cyber Resilience to Healthcare

Implementing cyber resilience for healthcare organizations delivers the following benefits:

- Increases your responsivity to remediate breaches faster
- Improves business continuity by providing enhanced capabilities for self-healing
- Reduces corporate risk and exposure by fixing issues as fast possible to minimize the loss of personally identifiable information for patients and clinical staff
- Provides a clear demonstration of due diligence to protect network assets and high value data

This topic is important because healthcare organizations must improve their network security. According to the 2021 Ponemon Institute Cost of a Breach Report, [2] healthcare organizations experienced the highest average cost of a data breach for the eleventh year in a row. The report also found that that average healthcare data breach costs increased from $7.13 million in 2020 to $9.23 million in 2021. This was a 29.5% increase. In addition, the report found that it took an average of 287 days to identify and contain a data breach.

While these data points are concerning, recent events are even more concerning. On February 23, 2022,[3] the American Hospital Association urged hospitals and health systems to remain vigilant for cyberattacks. The advisory states three concerns for the healthcare field:

- Hospitals and health systems being targeted by Russian-sponsored cyber actors
- Hospitals and health systems becoming collateral damage to Russian-deployed malware or destructive ransomware
- A cyberattack that could disrupt hospital's mission-critical service providers

While prevention of these potential threats needs to be a core area of focus, cyber resilience has a growing role within cybersecurity architectures. For instance, the NIST Cybersecurity Architecture Framework[4] relies upon cyber resilience activities as part of the Recover function within that five-point framework. The NIST framework is important for any government related healthcare facility or contractor because compliance to this NIST framework is mandatory.

---

[2] 2021 Ponemon Institute Cost of a Breach Report, Ponemon Institute, 2021. https://www.ibm.com/account/reg/us-en/signup?formid=urx-50915

[3] "U.S. Declares Start of Russia's Invasion of Ukraine, Introduces Sanctions; 'Cyber Shields Up,' Says CISA", American Hospital Association, February 23, 2022. https://www.aha.org/advisory/2022-02-23-us-declares-start-russias-invasion-ukraine-introduces-sanctions-cyber-shields

[4] NIST Cybersecurity Framework. https://www.nist.gov/industry-impacts/cybersecurity-framework

Creating a focus on cyber resilience is not new. In 2017, the Health Care Industry Cybersecurity Task Force also recommended that healthcare organizations focus on cyber resilience, according to this Health IT Security article.[5] Here are couple of the Task Force's specific recommendations:

- Increase the security and resilience of medical devices and health IT
- Identify mechanisms to protect R&D efforts and intellectual property from attacks or exposure

Where to Start With Cyber Resilience

Now that we understand what cyber resilience is, where should a healthcare IT team start? Again, this portion of IT security focuses on returning the IT network to a normal working state as soon as possible. This is accomplished by integrating visibility technology (like taps, external bypass switches, network packet brokers, and security monitoring tools) into your security architecture so that you can clearly see what is happening on your network and implement the proper responses.

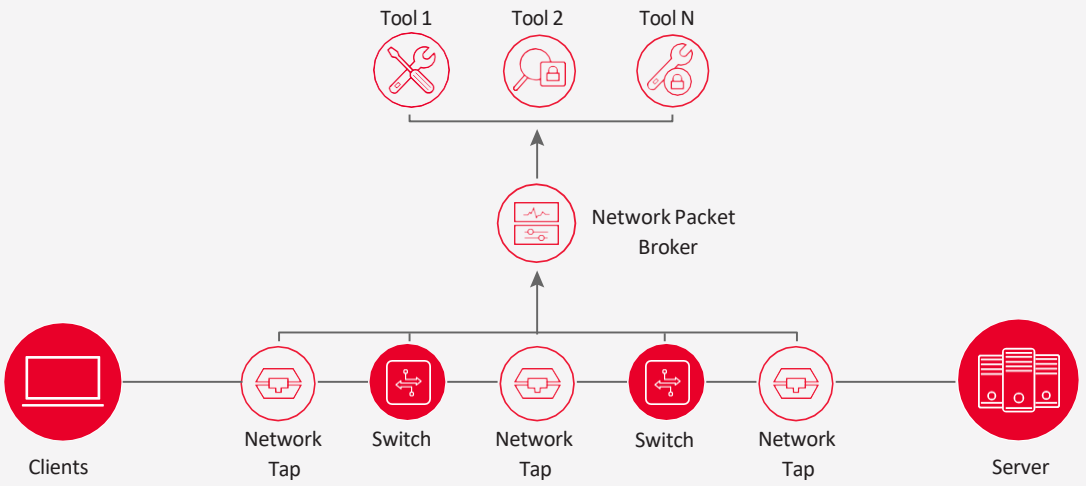

Figure 1. Example of a visibility architecture

[5] "Healthcare Cybersecurity Task Force Finds 6 Imperative Areas", Elizabeth Snell, Health IT Security, June 5, 2017. https://healthitsecurity.com/news/healthcare-cybersecurity-task-force-finds-6-imperative-areas

Specifically, this cyber resilience strategy can help in the following three general areas:

- Limit the amount of network downtime
- Decrease the time to discovery of an intrusion (or breach) — basically reduce the 6-month discovery time that the Ponemon Institute found earlier and reduce that interval to weeks or days
- Decrease the time from discovery to remediation of the intrusion (or breach)

The next sections will examine these three areas in detail.

Limiting the amount of network downtime

Here are five clear activities that you can implement as part of your network resilience strategy to reduce downtime:

1. Optimize network continuity by using external bypass switches with heartbeat messaging. These devices can be set to Fail Open or Fail Closed, as you choose.
2. Install inline NPBs to support fast architectural updates and reactions to security threats
3. Deploy inline and out-of-band network packet brokers using load balancing and n+1 survivability so you can maintain operations during “impaired” network situations. The right choice of packet broker can also provide reversion capability, which means that it can automatically sense when out of service security tools become operational again (i.e., if a security tool does a reboot and comes back online).
4. Install inline packet brokers with Active-Active processors to provide enhanced business continuity without loss of data. Active-Standby solutions will lose data while the standby processor comes online.
5. Use threat intelligence gateway to prevent exfiltration of data to known bad IP addresses. While a piece of malware may have “legally” entered the network, automatic updates (with new malware definitions) to these devices may catch and kill improper communications (like command and control, PII, and high value data) to off-network devices before those communications leave the network.

Let’s look at each area in more detail. For item number 1, an external bypass switch allows failsafe deployments of inline security and monitoring tools to ensure high availability and maximum uptime. The purpose of a bypass switch is to eliminate the pain of direct deployments of inline tools. While directly deploying inline security tools can create a line of defense, these tools can also result in single points of failure if connected serially. Even a strong mix of security and analytics tools can lead to network reliability risks as regular rebooting, maintenance, and upgrades of those tools will increase the chances of a costly network outage. In the event that an inline tool becomes unavailable, it can completely bring down the network link, significantly compromising network uptime and disrupting business continuity. This can be a significant problem for the almost 20% of IT enterprises that directly deploy inline security tools and the 40% that deploy internal bypass solutions instead of external-based solutions. In addition, when you replace various security tools, the integrated bypass may have to be removed as well. This causes an immediate outage and destroys any supposed bypass advantage. An external bypass

eliminates both issues. In addition, it also provides a better fail-over situation. Instead of multiple internal bypasses, you can have one external bypass switch. This is a simpler design with a solution that has more robust features like: heartbeat messaging to constantly detect trouble states, very fast fail-over times, and reversion capability to automatically reinstate traffic flows to security tools once they become available again. Once heartbeat messaging returns, the bypass functionality is disengaged, creating a self-healing solution.

Another technique is to install inline packet brokers. A good defensive architecture should have these already. However, if they haven't been installed, you will want them for a resilient approach. The packet broker gives you the ability to quickly add or replace security tools with minimal, to no, impact to the network. For instance, maybe a firewall or IPS is damaged (contaminated) by a piece of malware. If it's connected to a bypass switch or inline packet broker, then it can be rapidly replaced in the production network. The packet broker can also facilitate the addition of a honey pot to study ongoing attacks, load balancing to multiple concurrent tools, and the quick filtering of newly identified suspect traffic.

Network security and monitoring tool survivability is often thought about in terms of redundant tools, especially in the case of inline deployments. This is often called high availability or redundancy. In this situation, you have a duplicate copy of tools waiting in a standby mode to take over should the primary equipment fail. However, an option to high availability is to implement an n+1 solution for component redundancy. In this situation, you do not have a duplicate copy of tools. At the same time, you do not have to spend double the costs for a redundant solution like you do with high availability. You simply have one, or more, extra security tools that are fed by a network packet broker. In this solution, the traffic is spread across more tools than is actually needed and all tools are processing traffic simultaneously. Should one or more of the tools go into failure mode, the packet broker will redistribute the load across the remaining tools in the port group. Once the failed tool becomes available again, the packet broker redistributes the load. For example, let's say you need four IPS tools to process your inline network traffic. In this case, you would add a fifth IPS. The packet broker would then load balance the traffic across all five IPS tools. Should any one of the tools fail, the packet broker can load balance the full load across any of the remaining four IPSs. This provides a good level of survivability at a fraction of the cost of a fully redundant system. If you would like to have more survivability, like an n+2 situation, you can do that as well — all the way up to a fully redundant set of tools. It just depends upon the level of risk you feel comfortable with and your budget. This provides another "self-healing" component to your security architecture.

Depending upon your business needs, you may want another level of redundancy. Besides a redundant set of security tools, you may want equipment that has redundant microprocessors. However, there is an important distinction here as there are two options – Active-Active and Active-Passive. Active-Active means that both processors are working simultaneously to process traffic. Active-Passive means that only one of the processors is active while the second processor is in stand-by mode. Visibility solutions that are configured in active-passive mode will typically need a minute or more to restore full processing and restart data delivery. But a lot can happen in 60 seconds, and a lot of security issues can be missed. Packet brokers configured in active-active mode work with complete synchronicity to aggregate, filter, process, and deliver data to all security and monitoring solutions. This lets them work more efficiently, handle periodic traffic bursts, and failover in a few seconds or less to maintain continuous security inspection, without gaps.

The fifth activity mentioned above is to use a threat intelligence gateway. Threat intelligence gateways are typically a defensive approach to prevent infiltration, i.e., all incoming traffic from known bad IP

addresses are dropped at the edge of the network. However, if the IP address is not known as bad initially, packets (and malware) can infiltrate the network. If the threat is not captured by the IPS, it then continues into the network and does its dirty work. However, the threat intelligence gateway should have a constant update process with new access lists. This means that should malware get into the network; the threat intelligence gateway still has the possibility of catching outgoing traffic to newly discovered bad IP addresses and deleting that traffic before it leaves the network.

Decrease the time to discovery of an intrusion and/or breach

This section discusses another four activities that you can implement as part of your network resilience strategy to reduce the time to discovery (which helps the network to return to a normal state faster) of an intrusion or breach:

1. Start conducting cyber range training exercises so that you can recognize and respond to attacks faster. It's one thing to suspect that a certain type of attack has happened, or is happening, and another to be able to "see" the indicators of different types of attacks in real-time. Practice seeing these attacks in a cyber range is critically important. While you may not be able to tell a Petya attack from Ryuk, you can at least narrow down your search to the fact that it is probably a ransomware attack and proceed forward with that information.
2. Something else to consider is the deployment of out-of-band network packet brokers that support integration to SIEMs through a REST interface. This integration allows your network to support automation for faster data collection and the ability to thwart security attacks as fast as possible. Once a SIEM detects something suspicious, it can get the data it needs as fast possible to confirm or deny an intrusion.
3. Install an out-of-band NPB with SSL/TLS decryption. This allows you to decrypt once and distribute data in the clear to multiple analysis tools simultaneously. This improves tool performance (as the tool does not have to spend resources decrypting data) and it potentially improves time to resolution by decrypting once and passing that data to multiple security tools simultaneously.
4. The use of application intelligence can help find indicators of compromise (IOC). For instance, there may be the transfer of data off-network that can be spotted by looking at what applications are running in the network and where those application flows are located.

Cyber range training can be viewed as a resilient type of activity. While this is general training, rather than specific to resolving a single breach, this type of activity provides IT and security personnel with tips and tools for performing forensic activities. Basically, it gives you experience in recognizing different categories of attacks so that during, or after, a breach you can respond more quickly with the knowledge of what to look for and where to discover indicators of compromise.

Automation is another way to decrease time to resolution and increase resilience. A well-built packet broker should also be able to respond to commands from approved external devices through a REST interface. This allows you to automate responses to increase response times. For instance, maybe a SIEM sees suspicious traffic. With a RESTful connection to the packet broker, the SIEM can send commands to copy specific traffic (based upon IP address, VLAN, port tag, etc.) and send that traffic to an inspection tool (IDS, DLP, etc.) for further analysis. This automated capability delivers a phenomenal time to response to decrease the time to resolution.

Another thing to consider is centralized decryption for monitoring data. Network packet brokers deployed in an out-of-band use case can speed up problem isolation by using internal SSL/TLS decryption. This feature allows the packet broker to decrypt traffic and send it to one or more out-of-band monitoring tool, like a DLP, IDS, etc. Even if the traffic was decrypted and inspected initially as part of a defensive architecture and nothing was found, it could still contain malware, i.e., something got missed and it wasn't flagged or you didn't have time to inspect all of the flagged bad traffic, which led to an actual security breach. At this point, the traffic may need to be subjected to deep data inspection by an IDS or DLP. Even if the malware has already done its dirty work, you still need to understand how it works, where it came from, how it operates, etc. — which is information you can get from the unencrypted packet data.

Another activity you should consider is using application intelligence. For instance, one of the first things you want to do after a security attack is to find out whether your network has been compromised. One common approach is to buy a tool (SIEM, DLP, APT, etc.) that focuses on recognizing patterns and any "oddball" activity to uncover indicators of compromise (IOC). To do this, the tool needs network data fed to it by the packet broker. However, besides a basic data feed, the packet broker can deliver application-level data that can be used to see even more patterns. Basically, what applications are running on different portions of your network, how much data is flowing between different network segments, and geolocation of users. This helps you in one of two ways — you can either directly see IOC or the application data can be fed to your IOC tool.

Decrease the time of resolution of the intrusion and/or breach

A third category of cyber resilience functionality focuses on reducing time to resolution and remediation of the problem itself. Here are three more considerations for your cyber resilience strategy:

1. Use a security threat simulation tool with rich NetFlow data and then view that data in a lab environment to get a tactical analysis of how the breach took place
2. The ability to completely simulate the attack in your labs to validate any fixes is especially important. This is where you need a security threat generator, like BreakingPoint, to faithfully reproduce the security attack in your lab so that you can determine whether your security fix actually works. The last thing you want is to shoot yourself in the foot by rolling out a security fix that doesn't work. This could lead to another successful attack/breach and be a career limiting event for yourself.
3. Use a breach and attack simulation (BAS) tool to validate the fix in the live network immediately after installation of the fix. This allows you to test your network before a hacker gets a chance to test it.

An easy way to speed up the remediation process is to use replay and simulation tools. For instance, specific security threat simulation tools can also use rich NetFlow information from a packet broker to essentially create a capture of relevant NetFlow-based information made during an attack by the NPB. This allows you to get more specific data on the attack and how it took place. You can then replay that data in a lab to analyze the specifics of the threat and how it moves.

Closely related to this is the use of a security threat simulation tool to get a better look at how threats behave. For instance, once you have an idea of what you are looking for, you can run a simulation in a lab to see how the threat behaves. Specifically, let's say you think you are the victim of a Heartbleed

attack. You can use the simulator to accurately depict how that piece of malware will behave and then look for those tell-tale signs (specific registers compromised, protocols being used for messaging, etc.) on your network. You can also use this malware generator to validate any fixes in the lab. As mentioned earlier, the last thing you want is to miss something and create a self-inflicted wound from a supposed fix.

Lastly, once the fix is installed live in the network, you want to quickly use a BAS platform to run the same type of attack against your new network. The BAS platform will ensure that this is conducted in safe manner, of course. Once the network passes this test, you can inform your management team that the intrusion / breach has been fixed and that the network is back to normal operation.

Conclusion

Healthcare IT teams are under ever-increasing pressure to maintain a high level of network security and performance. To meet these challenges, they must implement security architectures that are both effective and maintainable. High touch solutions that require constant attention will not be sustainable — instead you need effective, self-healing architectures.

Cyber resilience is one of those techniques that enables the following:

- Limits the amount of network downtime from a security attack
- Decreases the time from intrusion to discovery of an incident
- Decreases the time from discovery to remediation of an incident

See the solutions available to you on Keysight's healthcare solution page.

Learn more at: www.keysight.com/us/en/industries/healthcare

For more information on Keysight Technologies' products, applications, or services, please contact your local Keysight office. The complete list is available at: www.keysight.com/find/contactus

KEYSIGHT TECHNOLOGIES