nixu
cybersecurity.

How companies get compromised?

The everyday life of cyber incidents and how to protect your business from them

WHITEPAPER

WHITEPAPER | May 2019

Contents

Executive summary

Targeted cyber-attacks have become a part of everyday life. Media reports on them almost daily basis. Data records are stolen or lost every minute. Companies worldwide are thinking: Have we been compromised? The real question for most cases is, how have we been compromised and why does it matter? How a successful cyber-attack affects the targeted organization varies, but based on recent examples the impact can be very significant.

How to protect your business then? There's a three-step model to handle cyber-attacks: 1. Detect, 2. Response, 3. Adapt. To protect your business successfully you must view security as a continuous process and have an active plan ready and available. Attackers are continually finding new ways to attack; modern technologies and service models such as cloud computing change the landscape and have created new opportunities for attackers. Because it is not feasible to build bullet-proof protection, you must also invest in your capabilities to detect and respond to security incidents. By continuously improving your abilities and adapting to the changing landscape, you will be able to protect your business-critical assets.

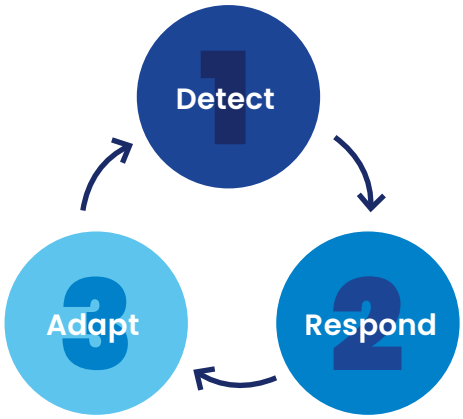

The first part of the continuous incident management circle is **detection.** Fast detection plays a crucial role in incident management, as it is the key factor in minimizing both data loss and the costs of incident response and recovery.

The second part of the continuous incident management circle is **response.** Timely and effective incident response will help you to get the situation under control before an attacker gets too deep into your systems.

The last part of the continuous incident management circle, **adaptation,** is often underestimated. When self-evaluation is done after every incident, the organization can improve its actions quickly and more efficiently. In this part, the protection, detection, and response are evaluated separately using a risk-based approach. Both a short- and long-term improvement plans should be crafted for continuous development.

1 Introduction

Data breach size and scope continues to increase from year to year. Meanwhile, targeted cyber-attacks continue to grow, enabled in part by the release of many exploits for critical vulnerabilities in widely-deployed software.

One alarming targeted cyber-attack was in headlines in March 2019, when the Norwegian aluminum company Norsk Hydro got attacked by ransomware called LockerGoga. The company had to shut down some of its plants, and it had to run its operations manually for several weeks. The company has declared it had around $50 million losses. However, the cybersecurity community has given praise for the aluminum giant of the way they handled the crisis, and luckily the company has robust cybersecurity insurance.

A more devastating attack was seen in December 2016, when the Ukrainian electricity grid was subjected to unprecedented cybercrime. A group of advanced hackers executed an attack simultaneously against several electricity companies, which resulted in approximately 225,000 customers being left without electricity in the middle of the coldest winter.

Based on data breach statistics it is likely that your company has already been a victim of some cyber-attack. It is also probable that you would not have noticed it. In many cases that Nixu has investigated recently, compromises were typically noticed by third parties rather than the organization itself. Often, the compromise had occurred a long time before it was noticed. An eye-opening statistic from the 2018 Ponemon Cost of a Data Breach Study revealed that on average it takes 197 days for an organization to detect a data breach.

The average cost of a data breach is $148 per lost or stolen data record. How many records does your company handle? One data record can be an email account, credentials to a system or personal ID number. In the UK, companies have experienced an average stock price decline of 5 % after a breach.

Should targeted attacks concern me?

Organizations sometimes believe that they do not have anything that government intelligence agencies want, and therefore targeted attacks are not something they need to worry about. With a perspective like this, the organization hasn't made a precise evaluation of their company security risks. It may not affect the everyday-life of your crucial production, but what about your liability as a company, your productivity downfall, or the third parties you're operating with? Nowadays companies use multiple sub-contractors, and usually, they are a sub-contractor to other companies. When you get hit by a breach, what will you say to your partner whose operations are down and they are losing money because of you? "We didn't think that cybersecurity would consider us"? Well, now you do know.

So what can I do to protect my business?

There is no single device or solution that would make you safe. If someone says that there is, they are wrong. Product vendors often say that their solution will make you resilient against all sorts of attacks. Unfortunately, this is not entirely true. The ability to protect your business comes from many different levels. In this paper, we have divided the actions you must take into three main parts: detect, respond and adapt. To successfully protect your business, you must cover all these areas. We conclude this whitepaper by looking ahead for the ever-evolving threat landscape and the next generation SOC.

2 How to detect attacks

When an attack occurs, the most important thing is time. A modern Security Operations Center (SOC) can detect any animalities found in the network and react to them in minutes at best. If SOC isn't in place, the reaction time can be months. Other important indicators are reconnaissance and logs. Take a quick tour on attack detection.

2.1 Indicators of being compromised

Typically, it takes time to understand that a company's ICT environment has been compromised. According to Verizon Data Breach Investigations Report 2019 that most breaches (36 %) take months to discover, sometimes it takes years (20 % of breaches). At the same time, the hacker needs only minutes to penetrate the environment and hide there. Fast detection plays a crucial role in incident management, as it is the crucial factor in minimizing both data loss and the costs of incident response and environment rebuilding.

In detection, it is important to understand what to look for. Indicators of Compromise (IOC) can be found from the protected environment using

a variety of methods. Here are eight indicators that a network security monitor staff is following:

1. **Bad traffic and anomalies** can be a sign of reconnaissance, penetration or call-home traffic. Bad traffic can be discovered with IDS/IPS or proxies, for example, and it includes connections to and from known botnets, port scanning, known exploits, and web application attacks. Many IP-reputation databases can be downloaded for free. Anomaly detection consists of the identification of unusual protocols and hosts, monitoring traffic baseline and rare queries from web servers. Calling home can be identified by the signature of the traffic or the destination of the connection.

2. **Internal network abnormalities** can be a sign of lateral movement. Getting visibility to what is happening inside your network helps when trying to find potential security breaches. Understanding the typical communication inside your environment is the key. Creation of watch-lists helps to baseline the environment.

3. **Connections or connection attempts** from one workstation to another may indicate that workstations have been compromised.

4. **Servers that initiate connections** to the internet or internal systems should also raise a flag: monitor newly opened services and new hosts connecting to servers.

5. **Changes in critical configurations,** user rights, and system errors should also be monitored. They can be a sign of persisting access. Configuration changes can be identified by a file integrity monitoring tool, network management tools and the system's audit logs. User right changes can be monitored primarily from user repositories and IDM systems but all the systems must be monitored for local accounts. System errors might indicate penetration or local privilege escalations. An attacker might also try to delete logs and typically it causes errors.

6. **Unusual user activity**. There are two reasons to monitor users: compromised accounts and malicious insiders. Both are equally important. The identification of anomalies is difficult without

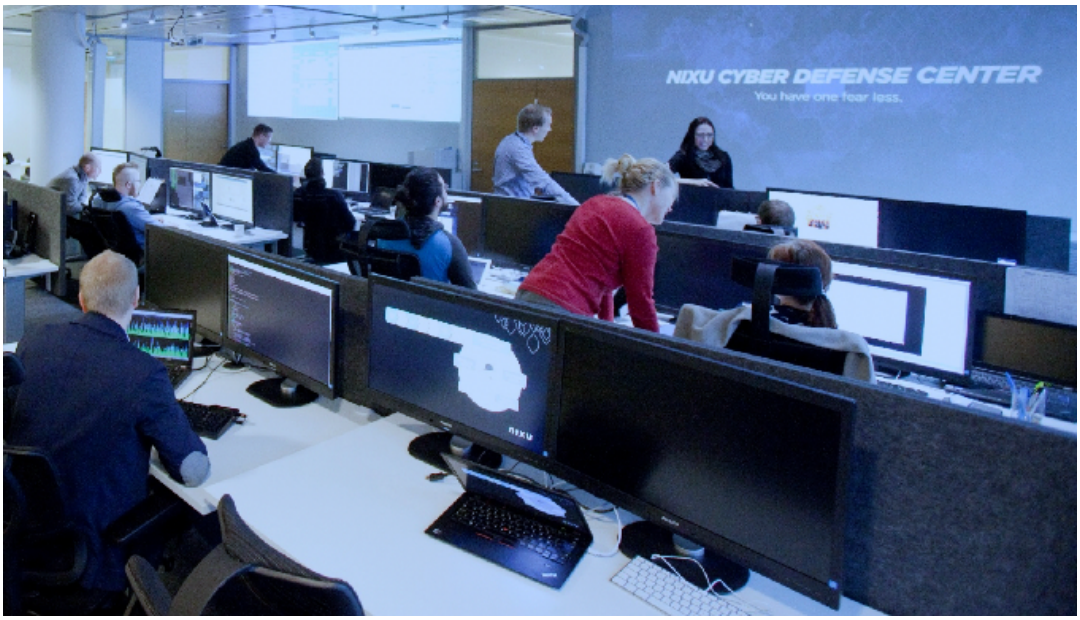

Picture: Nixu Cyber Defense Center. Fast detection plays crucial role in incident management.

proper tools such as advanced SIEM (Security Information and Event Management) or Identity and Access Intelligence tools, but it is not impossible. Even the most straightforward log management tools can monitor password changes, bad logins, sudo usage, and logins at unusual times and from unexpected locations. For example, remote login to a workstation with a local user account that should not be used in remote management could indicate that an account has been compromised.

7. **Unusual data behavior.** In a normal environment, it is data that must be protected – not the infrastructure. Understanding how data is used commonly and what unusual behavior looks like is essential. Network monitoring is not enough, because data can leak using a removable media. Tools to monitor data writing are audit logging tools and data loss prevention tools. Database Audit and Protection (DAP) tools can improve database activity logging.

8. **Data leaks and exfiltration in outbound traffic.** Data leaks are the most difficult to identify because all the confidential material must be tagged to be confidential if automatic tools are in use. Data leak monitoring is recommended for implementation based on the risk scenarios identified. Exfiltration is typically done in big chunks using standard file transfer methods. Those chunks can be identified based on the data size or the destination of the

transfer. Data can be exfiltrated in smaller pieces using DNS or ping tunneling or hiding the data in an HTTP request. Those slow-and-low types of exfiltration might be possible to detect, based on their malformed structure or the massive number of packets sent to a single destination.

Some of these indicators might be more crucial than others in maintaining your company's security posture. It's good to know your environment and its vulnerabilities. Current configuration management database (CMDB) with criticality and vulnerability data helps to understand the situation better. Audit and scan systems frequently to maintain this understanding.

2.2 Logs are the key

Logs play a crucial role in security monitoring. If managed correctly, they provide audit trails for troubleshooting and trustworthy forensics – and if not, the money is just wasted on the purchase of yet another ICT system. We at Nixu have come across many incident response cases where logging has been enabled, but the content of the logs was irrelevant to the investigation. From a distance, everything seemed to be just right.

The best practice in log management is to centralize log storing, alerting, reporting and retention to a separately managed server. Centralized logs server can secure the log data with proper access rights and ensure the integrity of the logs.

Log retention recommendations vary, but six months is a minimum and at least a year is recommended – but not all logs need to be saved for the same time. For example, in some cases, the audit trail must be stored for the audited data's lifetime. To make log storing simpler and more cost-effective, the logs should be stored as per the categories they represent: Changes in the access rights, critical data and critical security components of the environment should be stored in most cases for at least two years, whereas half a year could be enough for firewall access logs.

Logs provide audit trails for troubleshooting and trustworthy forensics.

It is also essential to collect the correct logs with the right amount of details. A safe choice is to collect everything from everywhere, but that

approach is expensive, capacity-wise, and slows searches. A better solution is to review log sources and estimate what information is needed to generate required outputs such as alerts and reports; as well as what is needed for a solid audit trail. Typically, changes in rights and configurations, user actions, errors, and alerts are needed at the minimum. All log entries should answer to following questions: Who? What? When? Where to? From where? Was the operation successful?

Log collection should be done quickly after they are generated and preferably using secure protocols. All the logs should have precise timestamps and log sources must have their clock set to the correct time and time zone. Network Time Protocol (NTP) is highly recommended for time synchronization.

A significant consideration in log management is workstation log collection. In typical environments, they are not collected, while in high-security environments they are. A good solution for that is to collect logs only from those workstations that are under special surveillance. That kind of workstations could be ones that have been suspected of having been compromised or ones where the user of the workstation is a more attractive target to attacks – as a company CEO. Prebuilt configurations (GPO for example) should exist and be tested so they can be quickly implemented when needed.

Log collection should be done quickly after logs are generated.

There are vast differences in log management systems, both in terms of price and functions. SIEM systems provide high visibility to the environment and the capability to correlate events from multiple sources. Those systems can be expensive, but if a company has data that is critical enough to protect, the investment will pay for itself. Alternatively, open source and freeware log management systems provide capabilities with a low cost of entry. Solutions at both ends offer better reporting capabilities than a typical ICT system.

Once the log management or SIEM system is acquired, its alert rules must be kept up to date. Detection capabilities of SIEM systems decrease quickly if the indication of compromise (IOC) list and lists of critical assets are not updated frequently. Besides, the log source integrations may not be valid after some time because the ICT environments tend to change so quickly.

3 How to respond to attacks

When a security incident is detected, incident response begins. Depending on how you have prepared for security incidents this may turn into a chaotic mess, or well-led and -organized work. The impacts of an incident can also be minimized with a well-led and -executed incident response. So how should an organization handle security incidents?

3.1 Initial response

After a cyber-attack or other security incident has been detected, there is a natural desire to start acting fast. This is usually a bad idea, as swift actions tend to cause more harm than good. The first and most important guideline is to stay calm and think before taking any steps. As stated earlier, it typically takes weeks or even months to detect a security incident. Typically, a few minutes or hours spent on planning your actions and response will not matter. On the other hand, well-planned and -organized efforts usually pay off.

Security Incident?

The first things that you should do are:

1. Contact your organization's security team or third-party incident response team

2. Collect what has happened, when and where
Facts about the incident are essential when planning response actions. How the event was noticed, and what was done when it was noticed?

3. Start making notes
Write all the actions you take, including the time they were taken. If you must use affected systems, write down precisely what you did. If you do not do this, it might be impossible to distinguish your actions from the actions of the perpetrator.

4. Find out the role and the owner of the affected systems
Incident response may come to the point when a decision to take systems offline must be made. To take this decision, it is important to understand the role of the system. Often taking down a critical system may cause more damage than keeping a potentially compromised system up and running.

> **TIP:** *There are also Incident Response Platform systems that have automated this process: they document the audit trail and alert the company stakeholders and show them their predefined tasks.*

Remember:

• Avoid using the infected systems if possible. Your actions may cause further damage.

• Never shut down the infected systems. Disconnecting the system from the network will cause network connections to die, which can make it impossible later to find out what the malware was communicating with.

• Never run antivirus checks or other extensive operations on the affected systems. They will ruin evidence and make investigation harder – including 'read-only' operations as they may modify filesystem metadata which could be used during a forensic investigation.

3.2 Preserving evidence

Before you take any actions that might change the state of potentially affected systems, you should decide your goal. Do you want to return systems back to production, or do you want to know what happened, how it happened, who did it, and how it was possible? If you choose the latter, then you must make sure that you preserve all the evidence before taking actions that might change the evidence.

The general rule of thumb is to start collecting evidence from the most volatile first and avoiding things that can change the system. If your security team is not able to carry out the memory dump quickly after the suspicion arises, then it is better to isolate the affected systems from the rest of the environment to avoid further compromises.

If you suspect that your workstation is infected, you should:

1. Take a memory dump to external media
2. Disconnect the system from the network or isolate it using firewalls or switches
3. Take a full disk image of the system

3.3 Identification

Before you can start resolving a security incident you must identify the extent of the it. What systems are infected, and what could potentially be infected? Typically this starts from the known facts of the event – what were the signs that made you suspect a security incident? These can be for example antivirus or SIEM alerts, a third party notifying you of suspicious traffic from your network, or a user noticing something out of the ordinary.

In targeted attacks, the patient zero, the place where the incident started, can be for example a laptop of a user who received a malicious document via email. Investigating this laptop will provide information about the malicious document used, other associated malware, potentially compromised user accounts, IP addresses and host names of command and control systems, and so on. With this information, you can start looking for other infected systems.

You may be tempted to find more about the IPs, URLs, and domains associated with the incident. You should always remember not to actively start communicating with systems that the attacker controls, i.e., do not ping, traceroute, or do DNS queries. These actions might tip off the attacker. SIEM solutions can help you find systems where potentially compromised credentials have been used, while firewall and IDS logs may reveal other systems that have been accessing the same websites where the infection comes from or communicating with the same command and control systems.

3.4 Containment, eradication and recovery

Before taking actions, you should consider the nature of the attack and whether your actions tip off the attacker. If the attack is a targeted one, especially if it is an APT attack, you should first monitor your environment to find out as much about the attack as possible. Only after you have a complete picture of the incident in hand should you start making coordinated actions to get rid of the attack.

Once you have identified the scope of the incident and you know what happened, you should start containment, eradication and finally recovery. Containing the incident means stopping it spreading, eradication means getting rid of the attacker and recovery means putting systems back into production. In this paper, we do not cover details about what could and should be done in each phase. Instead, we try to give a general idea about what kind of actions can be taken.

Containment methods vary depending on the incident. Cutting down the attacker's command and control channels, isolating affected systems and disabling potentially compromised accounts will usually stop the attack from spreading, but bear in mind that these actions might also trigger self-destruction mechanisms. This is why you should make sure that you have adequate backups and you have preserved the evidence without tipping off the attacker.

Preparation can help a lot in the containment. Having preconfigured firewall rules to block other traffic except what is mandatory for business makes it easy to block command and control channels. Having a predefined isolation VLAN ready makes it possible to isolate machines remotely.

Once the incident has been contained you need to start eradication. You should change the passwords of potentially compromised user

accounts. This must be done in a way that new passwords are delivered to users using methods that do not utilize potentially affected systems.

In many cases, it is easiest to reinstall infected workstations and systems. If reinstallation is not an option, make sure that you have identified all malware and other mechanisms the attacker used and remove them. When restoring data from backups make sure that backups are not infected, i.e., taken from the infected system.

After systems have been reinstalled or cleaned, you can start putting them back into production.

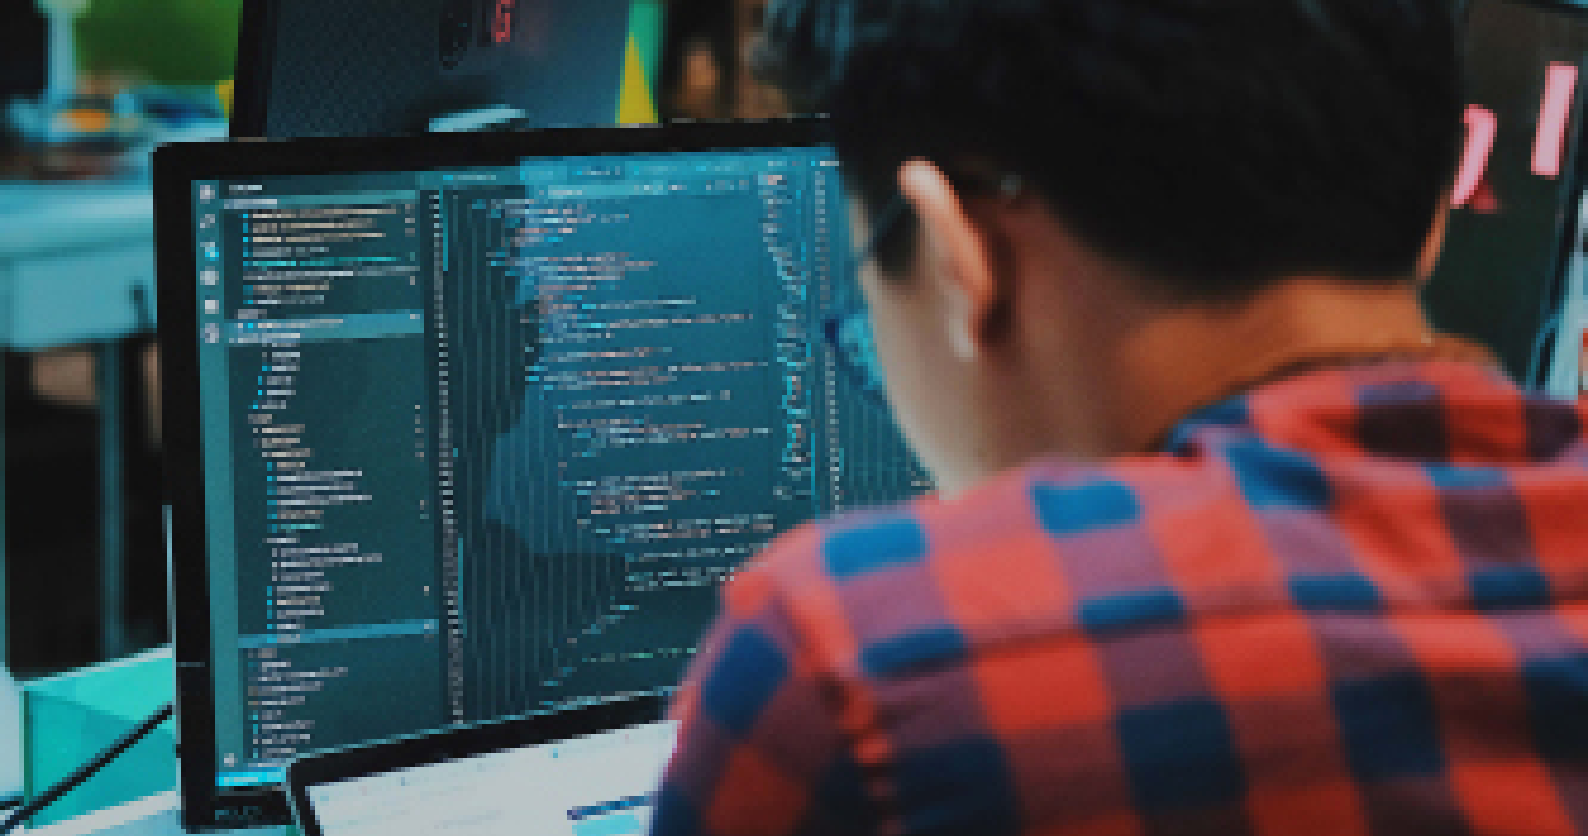

4 How to adapt after an attack

Once an incident has been resolved, it is vital that you take the time to analyze what happened and how you can improve your protection, detection, and incident response capabilities.

The first thing to do after the attack is to hold a lesson learned session where all the parties that took part in resolving the incident sit down and review what happened. The idea is not to judge the people involved but instead learn from the incident. What happened, how it was possible, how it was detected and how the incident was resolved. This usually reveals things that can be improved in the organization's security posture.

This is especially important as attackers are continuously finding new ways to bypass security controls and hack into organizations. Only this – often undervalued – self-evaluation can help the organization improve its actions quickly and efficiently. The key idea is to craft both short- and long-term improvement plans and try to continuously develop skills instead of starting massive improvement projects every four years.

4.1 After-incident monitoring

After the acute incident has been handled and the environment has been cleaned, it is important to continue monitoring of the environment; especially the systems that were infected. There is always a chance that some backdoor or another persistence mechanism from the attacker was not found during the incident response phase.

One good practice is to build lists of suspicious hosts and accounts for your monitoring systems. The idea is to have a lower alert threshold for these hosts and accounts. You will generate more false alarms, but as the alerts are only coming from a limited set of hosts and accounts, it is still feasible to check them. Lowering the alert threshold for the entire environment is probably not a good idea as there will be too many false alerts and you will not be able to investigate all of them.

4.2 Improving protection and detection

If the attack was not caught by the existing protection method, re-evaluation is needed. It is crucial to understand why the attack was successful and how you can protect against similar attacks in the future. Did we have a misconfiguration in our protection mechanisms? Weren't they up-to-date? Was the attack vector new? Are we developing our protection in the wrong direction? Do we have weak spots? Should we implement some new layers of protection? It is also important to remember that the choice should be made after a calculation of the cost (value) of the residual risks.

If the attack was not detected in time, something is wrong. When the protection is built based on risk estimations, it is possible that some spots of the environment's security are softer than others. It is crucial that you know what these weak spots are and that you at least monitor them to minimize the risk taken. It is very dangerous to make risk-based decisions with faint visibility of real events in the environment.

In addition to those questions asked in the evaluation of protection, the sensitivity of the detection sensors must be evaluated. Do we lack some critical sensors? Did we fail to get notification from existing ones? Or did we get too many false positives and it was impossible to identify the right alerts from that mass?

4.3 Improving incident response

Incident response concerns people and processes. Reviewing past incidents often reveals weak spots in procedures and instructions.

When an incident occurs, people are in a hurry and there is usually no time to start figuring out what to do. Processes created on paper that have not been tested in real life usually miss critical details about how people should work. And even if the processes are tested, nobody can test all possible scenarios. Therefore, it is important to take time after the incident and see if there is something in the process that could be improved.

Over time, your processes will be more fine-tuned, and you will be able to concentrate on the details. When you start, you are likely to face more significant issues.

For example, unclear security guidance may cause a situation where an incident is not reported to the security organization because people do not know how to report. A review of the past incident may also reveal technical weaknesses in incident response, such as an inability to make memory images from Linux systems, or a lack of competence to analyze the gathered evidence. As with processes, it is almost impossible to know all the tools you might need in advance. If you start seeing cases where a certain kind of tool or competence would help, you should consider investing in it.

5 Upcoming risks and the future of security operations

Targeted attacks are evolving, and as some threats are decreasing, new ones arise. For instance, cloud and mobile environments create a new landscape for attacks. However, the tools for detecting and hunting down anomalies are advancing.

5.1 Every cloud has a silver – and a cyber threat risk – lining

Cloud services are cost-effective and scalable and therefore a lucrative option for most businesses. Migrating environments and services to the cloud provide several benefits, but this transformation also includes risks.

When analyzing risks associated with cloud services, it should be considered that when the data used in cloud services is stored and processed in a service maintained by another party, the customer has very little control over how the cloud service is managed. In such a case, as many times witnessed, it is possible that data is misplaced, distorted, destroyed or disclosed to unknown third parties.

It is important that organizations carefully consider what data or services they intend to transfer to a cloud so that cloud services produce the best possible benefits and, also, organizations can keep risks under control. Furthermore, the security features of cloud services must be considered because the default settings of service providers may not offer the best protection there is.

5.2 Mobile security risks

As smartphones have become extensions of ourselves, company data has also found its way to new territories. Corporate data is now accessible from mobile devices, which creates new possibilities for cybercriminals. Data leakage, banking trojans, and phishing attacks are coming more relevant in the mobile environment, to name a few.

Modern mobile threat prevention can monitor mobile networks and prevent attacks from different devices, apps, and users. The same mantra is implemented in mobile environment than in desk computers: detect, respond, adapt.

5.3 Next generation detection tools

A modern security intelligence system works similarly to a security-oriented big data tool. It can be used to build a situational picture for different levels of the organization. The situational picture can include a view of the security status of an organization's core business structures, like core processes and resources.

The next generation of security intelligence systems provide context awareness and analyzing capabilities. Security intelligence tools, such as modern SIEM solutions, can bring business context to security events and in that way help in analyzing cases. A basic idea in security intelligence systems is that they gather a lot of data from an organization's environment and combine that with external knowledge in real time and with historical searches.

Collected data from an organization's own environment can be organization, asset, criticality, vulnerability, user, flow and event data. Human sensor data is also necessary because humans are sometimes targets of cyber-attacks and they have skills to detect phenomena. External knowledge and intelligence can be anonymous attack data from other organizations, research data from security labs, reputation data, behavioral data, warnings from authorities, and general threat

and vulnerability data. A concrete example of these is an automatically updated detection rule set or a reputation database in SIEM systems.

Another type of next generation tools for detection are Advanced Threat Defense (ATD) systems. Their signature-less approach for detection and protection goes far beyond traditional signature-based solutions. They are typically placed into three categories: host, payload, and network.

Host ATDs use agents in selected hosts that build baselines of the environment or the software that are typically run in hosts and identify anomalies in behavior. Payload ATDs are typically network devices that focus on testing the payload of the traffic in a sandbox and try to identify code that acts like malware or some other malicious code. Network ATDs search for anomalies in traffic with netflow or they record the whole of the network traffic for later forensic investigation.

ATD systems are not intended to replace signature-based detection, and the best results are achieved when those two are combined. These ATD systems may lack speed in protection and the possibility to be tuned to the maximum to achieve better performance.

5.4 Threat Hunting

A different approach to prevent cyber-attacks is not just to respond to anomalies and attacks but to hunt them down before they even hit you. The target of threat hunting is to search and find possible malicious codes and threat actors that have been able to bypass current prevention and detection capabilities.

Threat hunting is done by automation and humans together. Automation helps humans to gather and pre-analyze data, whereas humans are best to find other humans and rarely used attack techniques. Threat hunters in Nixu analyze existing data from Security Information and Event Management (SIEM) and End Point Detection and Response (EDR) solutions. They use Mitre Att&ck framework and Nixu's own Threat Hunting methodology. Threat Hunting produces reports that includes recommendations of what could be improved in the customer network environment.

Cyber threat hunting is an active way of improving organization's security posture. Instead of waiting for a warning or a potential threat, it's possible to detect them before they are realized.

Written by

Pietari Sarjakivi Director of Managed Security Services,
Nixu Corporation

Antti Nuopponen CISO & Head of Cyber Defense,
Nixu Corporation

Sources

[1] 2018 Ponemon Cost of Data Breach Study https://www.ibm.com/security/data-breach/
[2]2018_Global_Cost_of_a_Data_Breach_Report
[3]Ponemon Data Breach Impact Study UK 2019
[4]Verizon 2019 Data Breach Investigations Report p. 19
https://enterprise.verizon.com/resources/reports/2019-data-breach-investigations-report.pdf

nixu cybersecurity.

Contact

/nixuoy
@nixutigerteam
/company/nixu-oy

nixu.sales@nixu.com
www.nixu.com

Nixu is a cybersecurity services company on a mission to keep the digital society running. Our passion is to help organizations embrace digitalization securely. Partnering with our clients we provide practical solutions for ensuring business continuity, an easy access to digital services and data protection. We aim to provide the best workplace to our team of nearly 400 cybersecurity professionals with a hands-on attitude. With Nordic roots we serve enterprise clients worldwide. Nixu shares are listed on the Nasdaq Helsinki stock exchange.