ALIEN VAULT™

A PRACTITIONERS GUIDE TO ESTABLISHING A SECURITY OPERATION CENTER

This guide is intended to provide a technical audience the core information necessary to evaluate the security controls essential to establishing a Security Operation Center (SOC). This guide will provide a solid understanding of what data a SOC needs to be able to effectively operate and what methods can be used to gather that data.

Note: This guide is not intended to provide best practices or processes necessary to operate a Security Operation Center.

Why a Security Operation Center?

The world today is comprised of two types of organizations: those who have been breached, and those who do not yet know that they've been breached. The difference between these types of organizations is related to how soon the company can detect a breach and how effectively they can respond. An effective Security Operation Center provides the information necessary for organizations to efficiently detect threats and subsequently contain them. While eliminating the threats we face is an impossible goal, reducing the time it takes to respond and contain them is certainly achievable. With a focus on responding to, and containing threats, it is possible to define a Security Operation Center in simple terms: the technology and processes used to detect breaches and coordinate the appropriate response.

Historical precedent has lead most organizations to view Security Operation Centers as large, overhead-intensive, situation rooms staffed by huge numbers of highly trained (expensive) employees. This perception is due to the fact that the world has looked to the most security-advanced organizations (large financial institutions, telecommunications operators, and military organizations) as role models for establishing their security practices. While it is important to learn what we can from these institutions, we should ask ourselves, 'Do all organizations need to invest in the same way to detect breaches and coordinate responses? What are the threats my organization faces and how does that impact my investment priorities?"

Establishing a Security Operation Center is a necessary step for an organization to be able to detect and efficiently contain a breach. Once you've determined to establish a SOC, the next important question to ask is, 'how can my organization most efficiently achieve this goal?" To answer this, it is important to be critical of your organization's existing capabilities.

What technologies have you already deployed that might be of use?
Are these technologies available to my security team?
What capabilities do I have on my security team?
What are the current time constraints on that team?
What priority does the larger organization place on this project?

When establishing the Security Operation Center it is important that realistic understandings of these constraints are considered in order to ensure that an effective solution is created. If technology already exists, but access to the data cannot be guaranteed due to political reasons, it is of little use. Similarly, if technology is acquired but the overhead required for deployment, integration and management is beyond the capabilities of the current employees (either because of a lack of skills or a lack of time) then it will be of little help.

Essential Capabilities

The detection of breaches becomes an extraordinarily difficult task if the program is developed from the examples of a few sophisticated breaches. However, when considered in abstract the essential capabilities are quite basic. Securing an environment consists of answering a few basic questions:

- **What assets do I have to protect?**
- **Which of my assets are vulnerable to attack?**
- **How are people attacking my assets?**
- **How will I know if a breach has occurred?**
- **What actions are going to have the most impact on my security posture?**

Now, in a simple environment with only a few unsophisticated threats, this is not a difficult set of questions to answer. It is even possible to start evaluating these questions with manual processes and achieve a reasonable end result. However, it does not take much for the complexity of an environment to be beyond the scope of a manual process. For this, let us examine the technology that can be used to automate these tasks.

What Assets Do I Have To Protect?

What systems are critical to the ongoing function of your company?
Which systems are critical to the day-to-day tasks?
What other systems do those critical systems rely on?
Which systems manage and store sensitive information?

Having an understanding of the assets that your organization has is critical to prioritizing the efforts to respond to attacks and contain breaches. Being able to prioritize efforts to mitigate threats to a particularly critical system is essential for an effective Security Operation Center. Often, response policies and detection capabilities are broadly distributed without a good understanding of how they can be most effectively deployed with respect to business impact. A good understanding of the assets that are deployed and the services that those assets run is essential.

Beyond the ability to prioritize efforts based on potential business impact, having a good understanding of the assets deployed can help prioritize technical response to observed threats. Hackers often start without any understanding of the technical layout of the system they are targeting. As such, a full breach will often start with a discovery or probing phase. In this phase the attacker will scan the network and attempt to exploit known vulnerabilities on common services. Understanding what services are running on a particular host can allow observers of the attack to know whether it even has a chance of being successful; trying to exploit a Microsoft Windows vulnerability on a unix machine has no chance of succeeding.

Getting an accurate view of the assets deployed and the services running on those assets is a very difficult problem. The IT operations team can sometimes provide a reasonable picture in an asset management system. However, these systems are hard to access and are often out of date. An automated solution to this problem is preferable to ensure that the rapid rate of change found in modern data centers is accurately pictured in the security operation center.

To perform Asset Discovery, three automated approaches can be taken:

PASSIVE NETWORK MONITORING
By passively monitoring the network, traffic hosts and installed software packages are enumerated, identifying the protocols and ports used in the captured traffic.

ACTIVE NETWORK SCANNING
Active scanning probes the network to try and elicit responses from machines. Based on the response, the tool will identify the machine and the software installed on the machine.

HOST-BASED SOFTWARE INVENTORY
Installation of a host-based agent provides the last level of visibility. From an inventorying perspective the agent can enumerate all software installed on the machine, not just the software that is actively using the network (as required for passive network monitoring), or the software that listens on a port (as required for active scanning). This provides a far more comprehensive and accurate inventory.

These techniques can be employed individually or in concert with one another. Each approach requires a different amount of access to the environment to be inventoried; using a variety of approaches ensures that some information can be gathered even from tightly controlled environments. A good solution will provide a flexible approach to asset discovery, leveraging one or more of the techniques listed above to provide the most accurate picture possible. In addition, it will provide a centralized mechanism for discovering assets in remote segments of the network. Even a moderately complex network topology can cause operators to be faced with difficult access procedures, requiring manual steps to discover assets in remote network segments and adding to the long-term cost of managing the solution.

Which Of My Assets Are Vulnerable To Attack?

How are the services that I have running configured?
How can they be accessed?
Do any of them have known vulnerabilities that an attacker may be able to exploit?

Once you have a reasonable understanding of the assets deployed in your organization, it is important to understand where your weaknesses lie. When responding to an attack or breach, understanding how your organization may be exploited is a critical factor in prioritization. Knowing that a particular system is vulnerable to attack may make the environment that system operates in one of your top security priorities. For example, if a critical business process relies on an antiquated web-server that cannot be updated due to functional concerns, additional efforts can be made to ensure that the web-server is well protected from attacks. Or if a breach is detected on the same network segment as the web-server, additional priority can be given to the

incident response effort. It is not always feasible to remove a vulnerability in your system, but it is possible to understand where they lie and what impact might occur if they were to be exploited. Having a good understanding where your weaknesses lay is essential to understanding what efforts to prioritize when face with an attack or a breach.

Vulnerability assessment is a difficult and time-consuming effort. The software that is deployed in our environments, whether in the form of web applications, middleware or even the firmware in our devices, is all inherently vulnerable. Every piece of software has flaws that could potentially be exploited. This is largely due to the fact that security is often not considered during development. Even if security were a consideration, new methods for attack are always being developed, making the software that was secure a year ago vulnerable today. Because of the nature of this problem, vulnerability assessment is a continuous process, and ultimately a very difficult task to perform in an automated manner.

Understanding how a piece of software might be exploited often requires knowing how that software works, what types of data it accepts and what function it is supposed to perform. This is why the vulnerability assessment that is performed in our network environments simply tries to identify software packages that contain known vulnerabilities. Once the simplification of the scope of this task is done, automation can be performed.

The following approaches can be used to automate Vulnerability Assessment:

ACTIVE NETWORK SCANNING
An active network scan actively probes hosts using carefully crafted network traffic to elicit a response. This combination of the targeted traffic and the subsequent response allows an analysis engine to determine the configuration of the remote system and the software packages running on the system. This combined with a database of known vulnerabilities allows the analysis to produce a list of vulnerabilities that are present on the system.

HOST-BASED ASSESSMENT
Using access to the file system of a system, an analysis engine can perform a more accurate and comprehensive detection of vulnerabilities by inspecting the installed software and comparing the detected software packages with a list of known vulnerable software packages.

As both of these methods rely on a database of known vulnerabilities, it is important that they are performed periodically. Researchers publish information about new vulnerabilities constantly, and having an up to date database is the only way to ensure that your analysis engine can detect all of the latest vulnerabilities. As with asset discovery the deployment of vulnerability assessment can be a logistical hurdle. A good solution will provide a flexible approach in tightly controlled environments as well as provide for centralized management of the vulnerability assessment scans in environments with complex network topologies.

How Are People Attacking My Assets?

Is anyone attacking my systems?
What techniques are hackers employing when trying to compromise my systems?
Where are those attacks coming from?

Due to the nature of the Internet, there is not a single one of us who is not being attacked. Unlike the physical world, on the Internet it is quite simple for attackers to find us in an automated manner. The attacker may not know exactly whom they are attacking, but finding our front door is as simple as counting. Also unlike the physical world an automated attack against many systems is not any more expensive than an attack against a single system. As such, attackers are constantly scanning the Internet blindly attacking any and all systems they can find. For example, a simple web-server serving a web page for a private audience that had not been indexed by Google was faced with an average of 50 intrusion attempts every day. Every single one of us is constantly being attacked; understanding the nature, target, and sophistication of those attacks is critical when prioritizing our security efforts.

Threat detection can be thought of as the inverse problem of vulnerability assessment. Vulnerability assessment is the methodology for finding known vulnerabilities in your system; threat detection is the means to identify the attacks that are targeting those vulnerabilities.

Some attacks target particular vulnerabilities with known payloads; these can easily be detected using a signature that can identify those payloads. Other attacks are less well known and in these scenarios it can be possible to try and detect the technique the attacker is using to exploit the vulnerability. For example, when trying to exploit a buffer overflow vulnerability the attacker must inject the software with a large amount of data to overflow the memory segment the program is using for a particular buffer. Once this is achieved the attacker must then try to gain control of the execution of the program by writing machine instructions to the memory they have just overwritten. Since the attacker does not know the exact way their attack is going to be handled they must put in some room for error; this is often achieved by using a long string of 'no operation' instructions that will be essentially ignored by the machine. This, however, leaves a nice fingerprint for threat detection tools to try and find—a large amount of data being passed to a program that contains a long string of 'no operation' instructions can be identified as a potential exploit attempt. Alternatively, it is possible to detect indicators of an attack or indicators of compromise.

Threat Detection can be automated using the following approaches:

NETWORK INTRUSION DETECTION (IDS)
Analyzes the network traffic to detect signatures of known attacks and patterns that indicate malicious activity. This is used to identify attacks, malware, policy violations and port scans.

HOST-BASED INTRUSION DETECTION
Analyzes system behavior and configuration to identify behavior that could indicate compromise. This includes the ability to recognize common rootkits, to detect rogue processes, and detect modification to critical configuration files.

WIRELESS INTRUSION DETECTION
Accesses the wireless card to monitor wireless traffic and identify rogue networks. This allows for the detection of wireless clients, the associated networks and the encryption used. Critical for wireless policy enforcement.

As discussed before, targeted attacks often start with a discovery phase where the attacker is trying to identify vulnerable assets in your network. Detecting that a particular host is trying to probe your system can be essential to later evaluating whether or not a potential attack from the same host is real. The last resort to detecting threats is trying to detect indicators of compromise.

As with the other security capabilities discussed, it is essential that a flexible approach be taken in order to detect threats. Attackers explicitly craft attacks to evade different types of threat detection capabilities. Having layered, redundant, threat detection capabilities is the best way to ensure that an attack is identified.

Another important consideration is the placement of the threat detection capabilities that are deployed. Often these capabilities are only deployed on the perimeter of the organization with the thought that attacks only come from the outside. Modern attacks have broken this mold; today attackers leverage the fact that employees use their computers both inside and outside of the corporate firewall. A computer compromised while outside the perimeter now becomes a jumping off point for an attack from within the network. Threat detection should be deployed pervasively to address this. A good solution for threat detection will employ multiple techniques and provide substantial management capabilities to reduce the long term deployment costs.

How Will I Know If A Breach Has Occurred?

If I do not detect the attack how will I know an asset is compromised? If an asset is compromised how can I address it before the breach expands?

Not all breaches are avoidable. Our efforts to make an impenetrable system will never be enough to close all attack vectors. We will always be faced with a risk management decision when it comes to determining whether a vulnerability in an asset running a critical business process is worth the cost disrupting that process. Attackers know this and will always use this to their advantage. To ensure that the advantage they gain from this is as minimal as possible, it is important for us to detect a breach as quickly as possible.

Understanding the behavior of our systems and monitoring that behavior for indications that a breach may have occurred is essential to an efficient response. For example, if an internal server

that only sends and receives HTTP traffic suddenly opens up an outbound ssh tunnel to an external server, we have good reason to think that the server has been compromised. Understanding the behavior of your system as a whole includes understanding what assets communicate with one another, determining when services appear and disappear, and modeling network flow and protocol usage to detect anomalies.

Monitoring behavior can be performed with the following approaches:

ACTIVE SERVICE MONITORING
Actively validate that services running on hosts are continuously available. This is done with a network-level handshake and response providing feedback if the service becomes unavailable.

NETFLOW ANALYSIS
Analyzes the protocols and bandwidth used by each. This is done by capturing metadata from a TCP/IP stream, saving protocol information as well as calculating bandwidth usage.

NETWORK TRAFFIC CAPTURE
Captures the full TCP/IP stream. Allows for forensic storage of the stream so that detailed inspection can be performed if necessary.

HOST-BASED INTRUSION DETECTION
Can monitor the processes and resources used on a particular system. Detecting new processes or abnormal resource usage can be indicators of a compromise.

The information that can be gathered by behavioral monitoring tools must be used with caution. The systems run in our organization are far from predictable—seasonal peaks such as an end-of-the-quarter sales effort can cause loads and behaviors never seen before. A good solution for behavioral monitoring will provide multiple mechanisms for collecting this data. In addition it will provide a low-overhead mechanism for pervasive deployment in the organization.

What Actions Are Going To Have The Most Impact On My Security Posture?

What do I do first?
What data should I analyze today?
Should I stop a recently observed attack or try and contain a newly discovered breach?

When deployed at scale, the essential controls described to provide asset discovery, vulnerability assessment, threat detection, and behavioral monitoring produce a vast amount of data. The comprehension and prioritization of that data needs to be automated in order for decisions to made within a reasonable timeframe. In addition it is important for the data that is produced to be evaluated in conjunction with the data from the other security controls. Evaluation of each stream of data independently will lead to poor prioritization of efforts. For example, a new vulnerability is discovered on a host. Without having an understanding of what services that host provides or what attacks the host is being targeted with, it is hard to say whether it is more important to patch the vulnerability or to remove the malware just found on another server. However, if we know that the host in question provides our business partners with the ability to

refer customers to us and the server that has the malware infection is in our testing lab, then we can start to make better decisions. The ability to make sense of this data requires a system to consolidate and manage it all. That system must also provide normalization capabilities so that the data from disparate sources can be related to each other and the full picture can be presented to the ultimate consumer of the information.

There is only one approach to automating the comprehension of data today:

SECURITY INFORMATION & EVENT MANAGEMENT (SIEM)
An information management platform specifically designed for the evaluation of security information and events. Provides capabilities to normalize and analyze data from disparate sources and correlate it together to present a more complete picture of the incidents occurring in the overall system.

The most important aspect of a SIEM platform is how effectively the platform is able to correlate and present the data that is collected. The sole purpose of this platform is to increase the efficiency of the user who must analyze all of the data that is fed into it. Being able to associate the data together and respond to user query is necessary for a user to be able to understand all of the data, however it is not sufficient functionality to reach our goal of improving operator efficiency. The platform must be able to automate the correlation of the data in order to detect malicious behavior, large-scale attacks, and breaches. In a good solution this correlation is kept up to date with the latest threat intelligence allowing it to detect the newest methods of attackers and behaviors of malicious software.

Deployment Of The Essential Capabilities

Security organizations are faced with a difficult challenge—if a major breach occurs it will be their fault regardless of the support the organization has provided the security team before that date. The capabilities described above are essential for a security team to prevent breaches from occurring or containing breaches before they become a major problem. However, security teams are often given little political or logistical support within an organization. Business concerns such as budget, and resource allocation often take priority over proactive security. Given this limitation security teams must be as efficient as possible while establishing a Security Operations Center. In order to get a complete picture and comprehensive prioritization, the security tools providing the essential capabilities must be pervasively deployed throughout the organization.

To reduce the cost of this deployment it is important to try and find ways to use redundant infrastructure whenever possible. Most of the security tools require access to the far reaches of the network, whether to monitor the network traffic or to gather data about a host. If each tool is independently deployed then each tool will require configuration to understand your network topology and will require the proper credentials to gain access to every segment. In a similar fashion, a large number of these tools require rules or threat information to be periodically updated. If a common update mechanism can be used to provide rules to each of the tools then the amount of configuration will be greatly reduced.

Another consideration is the cost of integrating the individual security controls that are deployed. As discussed above, a SIEM product is designed to consume and evaluate the data that is produced by these products. However, a SIEM platform is only a means to an end. It is intended for the integration of any type of data and often requires substantial effort to integrate a new data source.

Conclusion

The capabilities discussed in this paper are the essential enablers for an organization to effectively and efficiently respond to threats and contain breaches. While the deployment of all of these capabilities is a daunting task, it is ultimately a cost- and time-saving project. Without having visibility into the actual threats facing your system, money will be inefficiently spent on compensating controls that might not have any impact. It is important for organizations to realistically evaluate how money is being spent and critically ask whether there is data from their own environment that substantiates the spend on these projects. Establishing a security operation center, the technology and processes used to detect breaches and coordinate the appropriate response, is the first step to long term, cost-effective, management of risk.

About AlienVault

AlienVault provides organizations of all types and sizes with unprecedented visibility across the entire security 'stack' with the AlienVault Unified Security Management™ (USM™) platform. Based on OSSIM—the de facto standard open source SIEM created by AlienVault—the USM platform has five essential security capabilities built-in: asset discovery, vulnerability assessment, threat detection, behavioral monitoring and security intelligence. The AlienVault Open Threat Exchange™, a system for sharing threat intelligence among OSSIM users and AlienVault customers, ensures USM always stays ahead of threats. AlienVault is a privately held company headquartered in Silicon Valley and backed by Kleiner Perkins Caufield & Byers, Sigma, Trident Capital and Adara Venture Partners. For more information visit www.AlienVault.com or follow us on Twitter.

040913

ALIEN VAULT™

www.alienvault.com