SANS
Analyst Program

Sponsored by
Bit9

Streamline Risk Management by Automating the SANS 20 Critical Security Controls

June 2012
A SANS Whitepaper

Written by: James Tarala
Advisor: G. Mark Hardy

Executive Summary

The 20 Critical Security Controls, a consensus project involving numerous U.S. government, private-sector and international groups, has received a great deal of attention recently as a framework of controls for defending organizations against cyber attacks. Today's cyber attacks are increasingly complex, advanced and able to remain hidden on the network for long periods of time.

As organizations begin to take the time to investigate and implement the 20 Critical Security Controls, they realize the document is more than simply a new list of things to do—it is a coordinated framework of controls that espouses a philosophy for combating the most common cyber attacks being observed today, namely those that blend malware, multiple attack vectors and manipulative social engineering tactics.

One of the core philosophies of these controls is that any defenses that can be automated should be automated. Rapidly detecting and thwarting follow-on attacks on internal enterprise networks minimizes an attacker's spread inside a compromised network.

Overreliance on manual assessment, response and mitigation has contributed to the current increased incidence of cybercrimes and compromised systems. In 2011, attackers breached more than 31 million individually identifiable records through hacking, malware, unintended disclosures and other similar incidents of cybertheft or cyber espionage, according to PrivacyRights.org.[1] PrivacyRights also reports that, in the first four months of 2012, the number of similar events had already reached 16 million.

What's worrisome is the persistent nature of these threats, which go undetected for months or even years. In an early case, intruders went undetected while compromising Heartland Payment Systems in 2008–2009.[2] The recent case involving the sophisticated Flame malware package found in Iranian nuclear control systems reportedly had been propagating, siphoning data and communicating to master controllers for more than two years before detection.[3]

Richard Clarke, the former U.S. government counterterrorism czar who reported to three different presidents, was quoted in *Smithsonian Magazine* as saying, "I'm about to say something that people think is an exaggeration, but I think the evidence is pretty strong. Every major company in the United States has already been penetrated by China."[4]

Whether this assertion is entirely true or not, it is clear that bad actors are conducting advanced, persistent cyber attacks against governments and businesses around the world. This paper discusses practical considerations for automating the 20 Critical Security Controls to create a more defensible network against these increasingly automated, persistent attacks.

---

1 www.privacyrights.org/data-breach

2 www.usatoday.com/money/perfi/credit/2009-01-20-heartland-credit-card-security-breach_N.htm

3 www.forbes.com/sites/johnvillasenor/2012/06/04/the-flame-cyber-espionage-attack-five-questions-we-should-ask

4 www.smithsonianmag.com/history-archaeology/Richard-Clarke-on-Who-Was-Behind-the-Stuxnet-Attack.html#ixzz1uCoQwdvg

Core Principles and Automation

In 2008, the Center for Strategic and International Studies (CSIS), in partnership with the SANS Institute, developed the 20 Critical Security Controls to address advanced security challenges faced by enterprises and governments. According to the guidelines, advanced attacks require improved, continuous security defense and response on the part of enterprises, as shown in Figure 1.
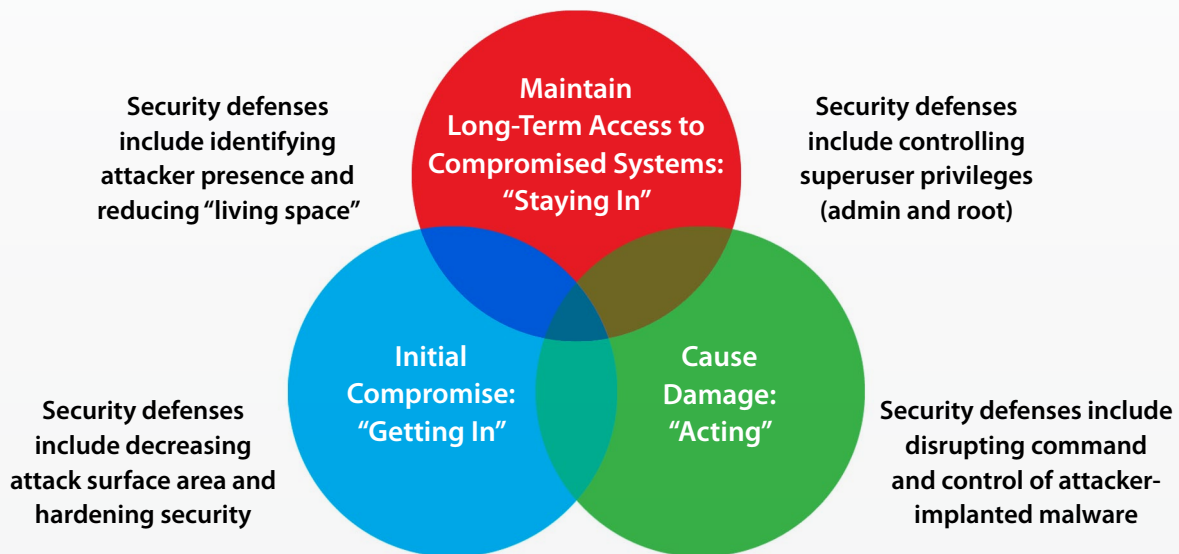
Security defenses include identifying attacker presence and reducing "living space"

Maintain Long-Term Access to Compromised Systems: "Staying In"

Security defenses include controlling superuser privileges (admin and root)

Initial Compromise: "Getting In"

Cause Damage: "Acting"

Security defenses include decreasing attack surface area and hardening security

Security defenses include disrupting command and control of attacker-implanted malware

Figure 1: SANS 20 Critical Controls Depiction of Advanced Attack Activities and Associated Defenses[5]

According to the project's website, the consensus document "… begins the process of establishing a prioritized baseline of information security measures and controls that can be applied across federal and commercial environments. The consensual effort that produced this document identifies 20 specific technical security controls effective in blocking currently known high-priority attacks as well as those attack types expected in the near future."[6]

The security priorities were established with the help of more than 200 government and civilian groups, heavily influenced by feedback from the U.S. National Security Agency, the Australian Defence Signals Directorate, the U.S. Department of Energy Nuclear Laboratories, the U.S. Department of State, the U.S. Department of Defense Cyber Crime Center and many other groups.

---

5 www.sans.org/critical-security-controls/guidelines.php

6 www.sans.org/critical-security-controls/guidelines.php

Core Principles and Automation (CONTINUED)

The focus was not simply to create a new list of things to do, but also to establish a new way of thinking so organizations could holistically and systemically defend themselves against these advanced attacks. Six principles guided the development of the control priorities, all of which are greatly aided by automation:

1. "Defenses should focus on addressing the most common and damaging attack activities occurring today and on those anticipated in the near future.

2. Enterprise environments must ensure that consistent controls are in place across the organization to effectively negate attacks.

3. Defenses should be automated where possible and periodically or continuously measured using automated measurement techniques where feasible.

4. A variety of specific technical activities should be undertaken to produce a more consistent defense against attacks that occur on a frequent basis against numerous organizations.

5. Root-cause problems must be fixed in order to ensure the prevention or timely detection of attacks.

6. Metrics should be established that facilitate common ground for measuring the effectiveness of security measures, providing a common language for executives, information technology specialists, auditors and security officials to communicate about risk within the organization."[7]

Guiding principle number 3 directly addresses automation: "Defenses should be automated where possible and periodically or continuously measured using automated measurement techniques where feasible."[8] Unfortunately, many organizations do not automate security measures. Instead, they rely on manual efforts to generate, aggregate and analyze valuable data about the way their information systems operate, including data sets describing network attack traffic.

A failure to automate the collection, aggregation and analysis of such data leads to increased costs and gaps in coverage. Man-hours cost more than computing time, particularly when you consider how much longer it takes humans to analyze incidents across the subnet or enterprise in which they occur. During this time lapse, events are not getting remediated, and defense against persistent threats and cyber attacks is not being improved.

---

7 www.sans.org/critical-security-controls/guidelines.php

8 www.sans.org/critical-security-controls/guidelines.php

SANS Analyst Program   3   Streamline Risk Management by Automating the SANS 20 Critical Security Controls

Vivek Kundra, then the U.S. government's federal Chief Information Officer (CIO), stated in a blog post that the security manpower and documentation for 150 major IT systems for the State Department alone cost $133 million in 2010.[9] In 2009, the White House reported that in federal agencies 60,000 Full Time Equivalent (FTE) positions were primarily responsible for security-related duties. With an average cost of $159,000 per FTE, the annual cost for these employees exceeded $10 billion.[10] The same report published the graph shown in Figure 2, illustrating the number of security-related FTEs by agency.


Figure 2: Number of Security-Related FTEs at Government Agencies in 2009
(Taken from www.whitehouse.gov/sites/default/files/omb/assets/egov_docs/FY09_FISMA.pdf)

---

9   www.whitehouse.gov/blog/2010/04/21/faster-smarter-cybersecurity

10   www.whitehouse.gov/sites/default/files/omb/assets/egov_docs/FY09_FISMA.pdf

Conversely, the cost to implement one information security sensor for an enterprise of less than 10,000 computer nodes is less than $100,000 annually, based on informal interviews with information security product vendors and enterprise CIOs.[11] For less than the cost of one FTE, an enterprise can implement one sensor that can automate some of the security for that enterprise. Most organizations will need more than one sensor to defend their systems, but with next-generation threat protection sensors that are much more accurate and centrally manageable, automated sensor deployment can cost significantly less than manual analysis.

The more pressing danger is the potential of more attacks succeeding—and persisting—due to a lack of automated controls. Next-generation threats are increasingly automated, targeting end users and organizations via web browsing, e-mail attachments, mobile devices and other vectors. Attackers set malicious code to operate covertly and automatically tamper with the system's contents, capture sensitive data and spread to other systems. Advanced malware avoids signature-based and behavioral detection through polymorphic coding techniques and can even initiate countermeasures such as disabling antivirus tools running on the targeted system. By always morphing their codebase, these threats stay ahead of signatures, blacklists and reputational updates. Therefore, they are unlikely to be blocked from stealing data and communicating it back to an attacker without the help of finely tuned, automated technologies.

Knowing that manual defenses are impractical and unrealistic, most organizations have begun the process of automating their monitoring systems. They have implemented network firewalls with automatic access control lists and intrusion detection sensors that automate the use of signatures. Also, next-generation threat protection tools have become more automated with multiple means of detection and analysis of advanced threats and zero-day attacks. Using virtual machine analysis and behavior-based anomaly detection, along with trust-based application whitelisting and other criteria, new security technologies can complement and enhance traditional signature-based detection, like Intrusion Detection Systems (IDSs).[12] Other automated protection and response processes include security to analyze all attachments entering the organization and to quarantine those that contain malicious code or file types, such as those that fall outside the organization's specified business rules.[13] Today's automated tools also include deep packet inspectors that can decrypt packets as needed, access and application monitoring systems, and proxy servers to automatically block access to unwanted network content.

Organizations should establish an incident response process that allows their IT support teams to supply their security teams with samples of malware running undetected on corporate systems.

In addition to their malware defenses, system vulnerability and configuration management, system inventories, application security, endpoint/mobility management and other similar information assurance controls can and are being automated. These systems collect data (sensors), aggregate and integrate data sets from multiple sensors (aggregators), and report and analyze the aggregated data (analysis engines). Let's consider each one of these roles in detail, including examining how they interact through further integration and automation.

---

11 Informal interviews were conducted by the author.

12 www.sans.org/critical-security-controls/control.php?id=5

13 www.sans.org/critical-security-controls/control.php?id=5

Enable Infrastructure Sensors

The first steps in any automation effort are defining what data needs to be collected and establishing effective collection practices. Organizations must be able to detect the most damaging attack activities occurring today, including advanced malware that bypasses traditional controls. Then, they must collect threat data in a way that can be automated, aggregated, analyzed and reported as useful information. To consolidate data and automate its collection, organizations must deploy tools and sensors to collect meaningful data from systems. The goal of a sensor is to collect data, not to analyze or aggregate the data (although many tools also provide these services).

Without the right sensors monitoring for multiple signs of breaches, organizations can't even analyze the data needed to protect themselves. If an organization does not know an advanced attack is underway, it cannot protect against it. If they don't know a particular risk exists, the exploit could go unnoticed for weeks or months.

Sensors can analyze all inbound and outbound data for signs of compromise in addition to collecting information that may be used to make appropriate and prioritized decisions regarding risk-mitigation efforts. There are 45 different types of sensors defined in the SANS 20 Critical Security Controls. Many of these sensors are being combined in advanced gateway technologies that include next-generation intrusion prevention, inbound and outbound web protection, application protections and more. Some of the sensor types include:

- Asset tracking systems

- Application whitelisting software

- Vulnerability management systems

- Patch management systems

- Antimalware software

- Network proxy servers

- Intrusion Detection Systems (IDSs)

- Authentication and access control systems

- File Integrity Assessment (FIA) systems

An example of a sensor would be a malware-analysis system. Ideally, organizations would deploy malware sensors that detect more than just signatures of known malicious code. Placing these signatureless sensors across communication channels (web, e-mail and file sharing) provides the capability to detect inbound attacks that use zero-day malicious code. They can also capture outbound callbacks to criminal servers, as well as abnormal conditions, such as unusual processes that are running on a system but were not authorized as a part of a standard system build.

By deploying network-based sensors on Internet and extranet DMZ networks, organizations can look for unusual attack mechanisms and detect network compromises. Also, a properly configured network-based threat protection sensor can automatically block bad traffic. Such confirmed signs of cyber attacks and processes could then send alerts to a central management system (discussed later in this article). Simply put, successful automation begins with quality data inputs, namely evidence of any abnormal condition that could be considered evidence of intrusion.

Not all the listed sensors need to be implemented before an organization can receive value from deployments. Even if only a few sensors are deployed, they have the potential to produce actionable results. For example, with access to only a vulnerability management system and a patch management system, an organization would be able to measure risk based on vulnerability severities defined by the Common Vulnerability Scoring System (CVSS). Security staff can use these scores to aggregate risk data about patches that were not installed on their systems, which would help them set priorities on which systems should be patched first in a patch-management cycle.

If an organization must prioritize its resources and start with a smaller set of sensors, it may want to consider beginning by implementing the controls defined by the Australian Defence Signals Directorate's (DSD) "Sweet Spot" controls. This research suggests deploying sensors on system management, access control and application whitelisting systems first when creating an implementation plan.

Aggregate Sensor Data

After data from sensors has been collected, it must be centrally consolidated and aggregated. Organizations must have the capability to take the data that has been collected by each of the individual sensors and automate the centralization, normalization and consolidation of that information to create a risk profile that key stakeholders can view and act upon. Armed with aggregated data, decision makers can make risk remediation decisions using the most complete set of intelligence available. It is equally important that this process be automated, because leaving it to humans to perform manually, even if done on a regular schedule, will result in gaps.

Too commonly, only a small, initiated group of individuals has access to separate data sets with little or no interaction between them. Investigators access only a small piece of the overall set of data needed to analyze risk and prioritize response. Other stakeholders, such as key business owners and senior executives, may never have the opportunity to observe a meta-view of the organization's risk. For example, an antimalware administrator is limited to information about malware infections, and a vulnerability management system administrator can see only vulnerabilities that are present at any given time on a system under his or her control. Those responsible for information security can make better decisions when data can be aggregated and made available to all decision makers.

During the collection process, the integrity of the data must be protected from compromise as well as from the data-consolidation techniques undertaken for the sake of efficiency. So in addition to normalized, consolidated data, raw data from sensors must be available to responders and auditors to ensure that they have the most accurate data possible. Sensors and aggregation tools must maintain the integrity of the data collected for these and other purposes.

Typically, this aggregation of data sets is performed by Security Information and Event Management (SIEM) systems. These systems are traditionally designed with the capability to interface with numerous sensors by default, thus saving the organization the time and effort necessary to create such interfaces themselves. SIEM systems should aggregate from what can be hundreds of disparate systems that provide log and event information. Infrastructure sensors ideally should analyze full packet headers and payloads of the traffic, including web, database and e-mail traffic destined for or passing through the network border. The resulting security, packet and log data should then feed into a properly configured SIEM system so events can be correlated against other network, user, application, log and system information feeding into the SIEM.

"SIEMs are like spreadsheets," said Alan Paller, the Director of Research for the SANS Institute on a public conference call announcing recent updates to the controls. His point was that SIEM systems, when related to the SANS 20 Critical Security Controls, are decision-support tools that assist decision makers by crunching through data sets and normalizing the data so analysts can search for patterns, behaviors and content that they understand. So, after collection, SIEMs must aggregate this data into one central system that enables decision makers to analyze all of the application, network, log, packet and security information they need during and after an event.

Analyze Security Data and Report Incidents

To help identify covert channels exfiltrating data through a firewall, threat protection can monitor outbound traffic for communications to known criminal servers, alert personnel about the source and destination addresses associated with these sessions, and deliver this information to the SIEM for automated actions, such as quarantining of the endpoint.

Effective threat protection requires strong analysis capabilities within the automated, integrated sensor/collection and reporting systems involved. Effective automation must extend beyond data collection and focus on infrastructure sensors to weed out the noise and false positives that IDSs and other monitoring/reporting systems are known to send. Further refinement must provide a way for responders and preventive personnel to perform meaningful analysis of the resulting data sets. Utilizing artificial or programmed intelligence, as well as virtual machines, can provide highly accurate data at reasonable speeds. Such devices should be able to support the analysis of information assurance data sets based on defined business analytics and metrics.

Alerts and the underlying logs that correlate with them must all feed into the analysis engine/SIEM to enable accurate results and provide a clearer picture of system security. Then the tool sets need to be able to provide the correlation information that today's advanced gateway sensors are able to collect, all the way down to deep packet information—and even data reconstruction—when, for example, exfiltration may be involved in the events. Figure 3 depicts how these controls and feeds would look in regard to Critical Control Number 1, Inventory of Authorized and Unauthorized Devices.


Figure 3: Sensors and Aggregation for Critical Control Number 1:
Inventory of Authorized and Unauthorized Devices

It's important to note there is a big a difference between basic log systems (aggregators) and SIEM systems. Some log aggregation systems, such as syslog or its variants, simply aggregate information into one place. By default, most log aggregators fail to provide any analysis of the underlying data sets. Log analysis engines also don't collect security information and alerts from network devices, vulnerability scanners or other security devices sending messages and alerts.

A SIEM is a "trainable" system that has the automated capacity for analyzing logs, sensor and other rich security information, including:

1. **Signatures.** SIEM tools should analyze the aggregated data sets for known attacks called *attack signatures*. Antimalware vendors have been using these signatures for years to alert organizations of potential malware outbreaks. Along with the many alerts they collect from antimalware and other sensors on the network, SIEM analysis engines need to be programmed with signatures of their own. This helps them coordinate the data sent to them from various devices to identify trends indicative of attack and make notification.

2. **Multivector Threat Capture.** Today's advanced monitoring and data collectors are collecting more security information than ever before. They can see beyond ports and destinations contained in packet headers and analyze and correlate threat activities across web, e-mail, and file-based attack vectors. For example, sandboxing suspect traffic for execution in near real-time effectively automates the process of confirming whether suspect traffic is actually zero-day malware targeting sensitive passwords, data being exfiltrated, or whether malware callbacks and remote commands are being performed. The data collectors and monitors send the data sets to the SIEM to enable a larger-scale view of the event across the network.

3. **Application Control and Endpoint Sensors.** Application control and whitelisting endpoint sensors track all files, process activity in real time and provide a live inventory of all executable content across all systems. When coupled with network sensors, such as Intrusion Detection/Prevention Systems (IDS/IPS) and firewalls, these systems provide audit data that enables earlier detection of threats, better filtering of noise and faster investigation and remediation times.

    These endpoint sensors automatically detect whether the file arrives on the endpoint and report an event only if the file is not approved. A network event highlighting a potential threat will be followed by an endpoint event of a new unapproved file only if: (a) the file actually lands on a system, and (b) there is no policy or rule that automatically approves the file. It is worth investigating such events because it can reduce the initial thousands of events to less than a few dozen actionable events.

Monitoring and SIEM tools must automatically alert key stakeholders to deviations like the preceding ones in language that both business owners and system administrators can understand. Alerts should be correlated and prioritized with chain-of-command decisions implemented into the workflow process. Decisions should not be solely in the hands of technical staff: Sharing more risk data with business owners increases the likelihood that appropriate resources are dedicated to responding to each identified risk.

These automation processes will still involve the use of humans to analyze the results, but it frees them to use the results for more than just response. It allows them to take actions to improve the overall risk posture. Automated intelligence provides the data humans need to make decisions and share what they know with each other, just as the criminals perpetrating advanced persistent attacks are doing. As more information on common attacks becomes available, more off-the-shelf products will provide automated analysis to enterprises.

Sharing attack intelligence among government organizations, private-sector enterprises and information assurance vendors will result in more automated intelligence being included by default in these products. To this end, the Department of Defense announced in May 2012 that it is expanding its Defense Industrial Base Cyber Security/Information Assurance Program (DIB CS/IA) to include providing classified threat and technical information to commercial businesses.[14] In addition, the U.S. House of Representatives passed the Cyber Intelligence Sharing and Protection Act (H.R. 3523) to facilitate the transfer of cybersecurity threat intelligence between the U.S. government and select private-sector companies.[15]

---

14 www.defense.gov/news/d20120511dib.pdf

15 www.govtrack.us/congress/bills/112/hr3523

Implementing the Controls

For organizations beginning their implementation of the 20 Critical Security Controls, it makes the most sense to start with sensors and tools already deployed and then determine what additional technologies/processes will be needed to fill out the automation program. The best place to start is by inventorying and updating existing technologies to meet defensive goals.

As an example, consider the U.S. Department of State's custom iPost system, which continuously monitors and reports risk on its IT infrastructure. According to the documentation, the iPost application consists of two main points of automation:

1. A single interface for administrators to access their monitoring data and objects

2. Single sign-on for enterprise-level management reporting across a variety of monitoring data[16]

As a result of collecting data sets from multiple sensors and aggregating that data into central sources, the iPost system allowed the U.S. Department of State to generate automated risk reports similar to the one shown in Figure 4.

---

16 www.state.gov/documents/organization/156865.pdf

Site_XYZ
Risk Score Advisor

Average Risk Score

| Site Risk Score	4,792.4
Hosts	63
Average Risk Score	76.1
Risk Level Grade	B
Rank in Enterprise	234
Rank in Region	27

The following grading scale is provided by Information Assurance and may be revised periodically.

| Average Risk Score		
At Least	Less Than	Grade
0.0	40.0	A+
40.0	75.0	A
75.0	110.0	B
110.0	180.0	C
180.0	280.0	D
280.0	400.0	F
400.0	-	F-

The Site_XYZ Site Risk Score was calculated as follows:

Component	Risk Score	Avg / Host	% of Score	How Component Is Calculated
Vulnerability	96.4	1.5	2.0%	From .1 for the lowest risk vulnerability to 10 for the highest risk vulnerability
Patch	807.0	12.8	16.8%	From 3 for each missing "Low" patch to 10 for each missing "Critical" patch
Security Compliance	1,089.0	17.3	22.7%	From .9 for each failed Application Log check to .43 for each failed Group Membership check
Anti-Virus	1,068.0	17.0	22.3%	6 per day for each signature file older than 6 days
SOE Compliance	975.0	15.5	20.3%	5 for each missing or incorrect version of an SOE component
AD Computers	3.0	0.0	0.1%	1 per day for each day the AD computer password age exceeds 35 days
AD Users	479.0	7.6	10.0%	1 per day for each account that does not require a smart-card and whose password age > 50, plus 5 additional if the password never expires
SMS Reporting	200.0	3.2	4.2%	100 + 10 per day for each host not reporting completely to SMS
Vulnerability Reporting	32.0	0.5	0.7%	After a host has no scans for 15 consecutives days, 5 + 1 per 7 additional days
Security Compliance Reporting	43.0	0.7	0.9%	After a host has no scans for 15 consecutives days, 5 + 1 per 15 additional days
Total Risk Score	**4,792.4**	**76.1**	**100.0%**	

Top 10 Risk Scores

Host_92898, Host_9270, Host_385336, Host_79133, Host_99015, Host_9271, Host_9275, Host_9237, Host_9283, Host_9276

(x-axis: 0, 50, 100, 150, 200, 250)

Risk Score History

(y-axis: 0, 40, 80, 120)
(x-axis: Mar 27 2009, Mar 30 2009, Apr 02 2009, Apr 05 2009, Apr 08 2009, Apr 11 2009, Apr 14 2009, Apr 17 2009)

For additional information on Risk Scoring, assistance with remediations, or to report suspected false positives, contact the IT Service Center to open a "Risk Score" ticket

Figure 4: An iPost System Automated Risk Report
(www.state.gov/documents/organization/156865.pdf)

At the beginning of their implementation of this customized risk reporting system, the agency had already deployed the following sensors in the organization:

- Security Center

- System Management Server

- Active Directory

Security and access systems certainly provide the most helpful information for correlation, analysis and actions to take. Many organizations, however, will not deploy a custom platform for this purpose. Instead, they will utilize existing SIEM systems to achieve this goal. Once quality assurance goals are defined, determine what existing tools can be leveraged and what new tools may be required by asking these two questions:

1. Are the correct sensors in place to collect enough security event and log data to determine good from bad network, access and system behavior?

2. Can the aggregation tool facilitate data from all needed sensors and provide the level of inspection needed to protect the network or investigate a breach?

If the answer to either of these questions is no, then establish a plan to remediate that issue—using existing systems wherever possible and supplementing wherever necessary to achieve those goals and keep implementation costs down. For example, organizations may need to upgrade their sensors from traditional firewalls and IDSs to include sensors that can open and inspect packets, perform application security functions, execute suspicious attachments in virtual environments and more, in what are coming to be called advanced security gateways. These gateways provide richer information for the correlation engine. Capacity should be considered in any expansion plan, because the SIEM/correlation engine must be able to support logs and events from any new sensors that are added.

Making Automation a Priority

Specifically, what steps should organizations take to make automation a priority? Take the following actions to begin prioritizing automated defensive systems:

1. Establish a baseline of sensors based on what is in your organization and the controls and sensors listed in the 20 Critical Security Controls to collect front-line data for defense.

2. Determine which of the sensors listed in this baseline have already been effectively implemented in the organization to collect useful defense data automatically.

3. Perform a gap analysis between the desired control, the sensor baseline and the system currently implemented in the organization.

4. Implement automated systems to address the gaps identified.

5. Ensure all data sets from automated systems are aggregated into central repositories for analysis.

6. Implement analysis engines to automatically analyze the centrally aggregated data sets for necessary remediation actions. Include reporting capabilities that can be shared with enterprise defenders and business owners.

7. Implement a continuous process improvement plan that encourages the collection of better data sets from sensors, ensures the aggregation of all useful sensor data and fine-tunes the analysis of the data collected based on emerging threats.

Conclusion

Manual information system defense is risky behavior. While the human eye will always be needed in protection and remediation processes, over-reliance on manual activities and security functions acting in silos results in longer windows of vulnerability. It is also costly: Government organizations have reported spending hundreds of millions of dollars on security personnel, and they still have experienced significant breaches of an increasingly advanced nature.

Automation, as much as possible, is strongly recommended in the SANS 20 Critical Security Controls. Although humans will always be needed to run and manage security operations, automation will help keep up with current threats and improve their overall risk posture.

Elements of automation include the use of increasingly advanced sensors to detect threats, automatically collect malware forensics, produce remediation details and alert to anomalous and malicious events. Collection, correlation, analysis and reporting/alerting of event data will need to continually improve as threats grow more advanced and persistent.

Continued automation, as well as the continued watchful eye of human analysts, must perpetually work together to protect against and remediate incidents in an increasingly sophisticated threat landscape. Automation should help beat advanced threats at their own game. It can be used to baseline security posture, determine what is normal, detect what is not and, ultimately, improve the overall organizational security posture through ongoing remediation and systemic improvements.

About the Author

James Tarala is a principal consultant with Enclave Security and is based out of Venice, Fla. He is a regular speaker and senior instructor with the SANS Institute, as well as a courseware author and editor for many of their auditing and security courses. As a consultant he has spent the past few years architecting large enterprise IT security and infrastructure architectures, specifically working with many Microsoft-based, directory services, e-mail, terminal services and wireless technologies. He has also spent a large amount of time consulting with organizations to assist them in their security management, operational practices and regulatory compliance issues. He often performs independent security audits and assists internal audit groups to develop their internal audit programs. James completed his undergraduate studies at Philadelphia Biblical University and his graduate work at the University of Maryland. He holds numerous professional certifications.

SANS would like to thank its sponsor: