



Interested in learning
more about security?

SANS Institute InfoSec Reading Room

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

Own Your Network with Continuous Monitoring

Copyright SANS Institute
Author Retains Full Rights



*Sponsored by
Tripwire*

Own Your Network with Continuous Monitoring

September 2012
A SANS Whitepaper

Written by: Jerry Shenk

What IS Continuous Monitoring? *PAGE 2*

Discovery *PAGE 3*

Analysis *PAGE 6*

Tuning *PAGE 8*

Reporting *PAGE 10*

It's a Big Job: Where Do You Start? *PAGE 11*

Introduction

Continuous monitoring is an ancient concept dating back to warring factions using arrows, clubs and spears. The Babylonians in 539 BC didn't think they needed to monitor their defenses because their defenses were so impenetrable—that is, until the Persians¹ dammed up the river to sneak in through what turned out to be an unmonitored vulnerability. More recently, we've seen references to multiple break-ins that relied on gaining a foothold through one or more vulnerabilities that may or may not have been known.

Because of continuous changes in the threat and monitoring landscape, over the past few years, monitoring has become so important that federal agencies are now required to continuously monitor their systems and defenses. Outside the federal government, IT organizations in almost every sector are required to maintain and monitor their computers to various degrees.

We know how to monitor firewall logs, Intrusion Detection Systems (IDS) and event logs. It's not easy, but at least we understand the idea and the data that these devices generate. But what else should we be monitoring? The United States Department of Commerce has released the Federal Information Security Management Act (FISMA), the credit card industry has its own standard, and computer security groups have the 20 Critical Security Controls. Unfortunately sometimes an overabundance of guidance can just add to the confusion—as is the case with continuous monitoring.

Lack of monitoring is why most breaches aren't detected for weeks, or even months, according to the 2012 Verizon Data Breach Investigations Report.² But, the process through which those breaches were detected can hardly be called monitoring because more than half of those breaches were discovered by external parties—and not by the breached organization that has been continuously monitoring.

Obviously, the goal is to do much better than that. In this paper, we look at what continuous monitoring is and how organizations can devise a solution that works for them.

1 Herodotus, Histories 1.190-191

2 www.verizonbusiness.com/resources/reports/rp_data-breach-investigations-report-2012_en_xg.pdf

What IS Continuous Monitoring?

Continuous monitoring is a cycle consisting of four basic phases: discovery, analysis, tuning and reporting (see Figure 1). Each of these basic phases has multiple parts, but simplifying the basic phases makes the entire process applicable to a wider range of situations. These are not individual phases that run in sequence; all four phases need to be going on continuously.

While many organizations talk about real-time detection and remediation, that generally isn't a realistic initial goal. By prioritizing and taking steps to improve, you should be able to detect deficiencies within 24 hours, with an eventual goal being detection in two minutes and resolution or isolation within five minutes.

Continuous monitoring processes will benefit from automation of security processes and workflow. To be successful, a continuous monitoring program must do the following:

- Automate aspects of the continuous monitoring cycle.
- Manage and maintain consistent configurations for the computing systems in the organization.
- Have management support to authorize active participation by IT personnel to maintain the approved configurations.
- Have management participation in identifying and approving the key assets that should be monitored.
- Provide reports to appropriate personnel.



Figure 1. The Continuous Monitoring Cycle

The 20 Critical Security Controls (20CSC, also referred to as the Consensus Audit Guideline or CAG20) state that through such controls it is possible to rapidly improve security. The guidelines go on to state the strategic advantage of such a program.

To make rapid progress, according to the introduction to the guidelines, the chief information officers (CIOs) and chief security officers (CSOs) of an organization need to “establish a prioritized baseline of information security measures and controls that can be continuously monitored using automated mechanisms.”³

The 20CSC (Twenty Critical Security Controls for Effective Cyber Defense: Consensus Audit Guidelines)⁴ is a consensus-based document that is useful for governmental and commercial environments. The 20CSC document would be a valuable tool to assist with any continuous monitoring initiative.

So, how do you get started? Work through each of the phases of the cycle to develop a plan that works for you.

³ www.sans.org/critical-security-controls/guidelines.php, Introduction

⁴ www.sans.org/critical-security-controls/guidelines.php

Discovery

The first step in a continuous monitoring implementation is discovery because, unfortunately, organizations often don't know what systems and versions they have on their networks. When hired, penetration testers often report vulnerable services to an organization that didn't even realize those services were available. Some of the reasons for these "surprise services" are the relentless growth of IT services, legacy hardware and software, test environments and a general failure to follow procedures and document changes. It all sounds so simple, but in reality, IT organizations are often expected to provide their organizations with new features on short notice with little concern for security implications.

Own Your Assets: Inventory Systems

The logical starting point in the discovery phase is to document and maintain an inventory of network assets including hardware, software and services. This inventory should include internal and external assets. Maintaining a valid inventory is a constant requirement, because most networks are constantly in flux. Hardware changes, software changes—and even the network infrastructure changes frequently. Once an organization knows what its assets are, it can prioritize the assets by value to both the organization and to an attacker.

The discovery phase of the continuous monitoring cycle is critical to success. Too many organizations discover new vulnerabilities only after they have been exploited. To avoid this situation, organizations need to constantly look inside and outside their systems for vulnerabilities and threats. The vulnerabilities are security issues related to the organization that could be exploited. The threats are agents that could cause harm to an organization. In our Babylonian example, the river that supplied water to the city was a vulnerability, and an unknown vulnerability. The Persian army was the threat. IT people are aware of many vulnerabilities; but, often the worst vulnerability is the unknown vulnerability.

The asset discovery phase should consist of documenting all the known computing devices in the organization, including test equipment, spare equipment and any offline labs or devices. It may not be possible to monitor offline computing assets as effectively as online equipment. Still, such items should be included in the inventory so that when updates are required, that equipment is not missed. Test labs and equipment are similarly important for inventory monitoring because personnel often bring test servers online to perform tests, apply patches or be used for some other short-term purpose and leave them online past the expected time frame. When this happens, those devices become security liabilities to the network.

The asset discovery phase should also include passive traffic monitoring. Passive monitoring will allow an organization to detect devices on the network that are not expected. Passive monitoring can also be configured to detect traffic patterns that are not expected. One of the starting points for passive monitoring can be the detection of traffic with known or suspected hostile networks including botnets, unfriendly countries and networks that other organizations have reported. Place these types of passive monitoring at network ingress and egress points.

Know Your Aggressors: Assess Threats

Discovery also needs to include threat agents—those who might want to steal data of value and/or inflict harm on an organization. As you consider the threat landscape, you would be wise to include and prepare for events that could cause system outages but are not necessarily IT security events, for example, power interruptions, natural disasters, fires or floods. The NIST Special Publication 800-37 (Guide for Applying the Risk Management Framework to Federal Information Systems)⁵ recommends that risk management involve every aspect of the entire organization.

Continuous monitoring from an information systems perspective is referred to as Tier 3 on page 7 of the NIST publication. This is what the U.S military calls *situational awareness*. It goes beyond knowing what's on your network, what services you have enabled and where your most valuable assets are—to include knowledge about the enemy. If the enemy just developed an exploit for an area that you think is secure, you need to test it and find out before it is used against your organization.

Along with natural events that cause system outages, threats stem from two types of threat actors:

- Privileged insiders, either through accident or on purpose, leveraged in the Stuxnet breach in 2010 and the Aramco breach reported in August 2012
- External threat actors motivated by politics, greed, activism or even sponsored by governments

Determine your potential attackers' motivations by figuring out what they want (your assets). It is also vital that you recognize the types of harm these different groups might try to inflict on your systems.

To understand the threat landscape, threat actors and their motivations, follow their new exploits by monitoring security newsgroups or security-focused websites such as the Internet Storm Center,⁶ CERT,⁷ or other sites that digest the security news to provide the most critical information.

To stay ahead of the attackers, use these outside resources, and compare what you learn with what you know about your assets and their state. If you were the enemy, how would you find and exploit weaknesses among your systems and their users/operators? Don't get worried about whether it's possible or not, that's the attacker's problem. Then, try to make the attacker's life more difficult by layering in additional defenses, setting up more stringent logging practices in vulnerable areas and using custom monitoring on your most critical resources.

5 <http://csrc.nist.gov/publications/nistpubs/800-37-rev1/sp800-37-rev1-final.pdf>

6 <http://isc.sans.edu>

7 <http://www.us-cert.gov>

Know Your Configurations: Audits and Auditing

Critical Control 3 deals with defining and maintaining secure configurations for laptops, workstations and servers, and Critical Control 10 does the same for network devices and firewalls. Organizations that do not have a plan for achieving and maintaining good configurations for their equipment fall into the trap of rolling out equipment and software whenever it is required, rather than using a controlled process. Multiple internal groups also end up with different software to do similar jobs, which quickly becomes difficult, time consuming and expensive to manage. Even worse, when a security concern comes to light, it is difficult for the organization to know which devices need the fix and which ones don't, because they have no idea of the actual configuration of the machine(s) in question.

As a rule, default passwords and usernames should be changed. If remote management options are enabled, consider disabling them or restricting access. Attackers don't have to work hard to exploit systems still using default settings and services. Systems using default settings are often referred to as "low hanging fruit." Attackers and code builders simply look these settings up on the Internet, application support sites or hacker sites, or they purchase a target system and do their own discovery.

By maintaining documentation on how software and hardware should be configured, an organization removes this low hanging fruit and makes it more difficult for attackers to unleash their malware packages. Once a consistent configuration has been determined, continuous monitoring against this baseline configuration can verify the system state. For example, during an event, knowing your systems and their configurations can help testers determine if your specific builds are vulnerable to newly discovered threats.

Analysis

Once an organization knows what assets it has, the next step is to analyze its systems to see what threats and vulnerabilities are relevant to the organization. Every open port enables a possible connection from an enemy. Every permission granted to a user could be exploited by the user or by an external party who gains control of the user's system. While business reasons drive most equipment and software installations, it is important to review each new implementation to make it as secure as possible.

Own Your Vulnerabilities: Network and System Assessments

From a continuous monitoring standpoint, a vulnerability assessment is an automated process that periodically scans the network for services and hosts in an unexpected state and, therefore, vulnerable. An internal vulnerability assessment should also be able to authenticate to servers and workstations to verify software versions, revision levels and that expected patches have been applied.

This does not mean just looking at IT systems, either. What about the systems that support critical network devices? Imagine the effect of installation of the recently promoted power strip⁸ that was hacked to hide a computer that could connect to a server outside the network over an encrypted tunnel.

It takes imagination to think of all the possible vulnerabilities. Only when you have thought of the vulnerabilities can they be analyzed. Combine knowledge of the discovered vulnerabilities with knowledge of the threat scenarios mentioned earlier in this section, and organizations can be aware, work around the vulnerabilities and provide protective controls before an attack occurs.

It is also possible that a service you think is safe today could, in the future, be identified as being vulnerable. Cataloging the state of all software ahead of any incident can make it easier to resolve such problems as they unfold—rather than after they've done damage. To prevent problems you don't know about yet, avoid installing software and services you don't need. As with most aspects of continuous monitoring, a team approach to identifying vulnerabilities and threat vectors can be helpful by providing different perspectives.

Own Your Network: Traffic Analysis

Continuous monitoring also means watching your inbound and outbound traffic. This is most commonly handled with an Intrusion Detection System (IDS), which can monitor the network traffic and identify traffic to/from known problem networks, such as botnets and other network blocks that have a reputation of supporting malicious activity. They can also identify anomalous behavior you should analyze to determine whether the traffic poses a real threat to the systems and configurations in their environment.

The traffic analysis system should probably be the same network monitoring system discussed earlier under network discovery. You should already know what IP addresses are in use on your network. If you see traffic from hosts that you thought were not in use, track down the source and either update your documentation or shut down the unexpected hosts.

⁸ <http://pwnieexpress.com/products/power-pwn>

Know Your Systems: Analyze Logs

Even the most precise attacks involve crashing services, logging users in at times that are not normal for them and/or transferring files that are not normally transferred. All these activities could be detected from log analysis.

Don't try to go from zero to 100 in a moment. Start by collecting the logs centrally and verifying that auditing is set up correctly for the hardware and software platforms supported by your organization. *Important: Verify that you have enough storage space to collect the logs.*

Once you are collecting the logs, start running some simple nightly reports on the most critical issues. Generate some baselines—a simple report that counts the number of successful and failed logins is a good place to begin. If you see a spike in failed logins, that's a sign that you should look deeper. Looking at ingress firewall logs can give you an idea of whether your network is “under surveillance” by an unknown entity. Looking at egress firewall logs with the same baselining approach can give you an idea of whether some internal workstation is attempting to “phone home” to a master controller. Depending on your firewall and your network activity level, you may be able to log each inbound and outbound connection, along with the duration and amount of data transferred. Monitoring for large file transfers can help identify data exports. For additional “Quick Wins,” review the 20CSC paper. Be sure to review Critical Control 14, which relates to log analysis.⁹

⁹ www.sans.org/critical-security-controls/control.php?id=14

Tuning

Continuous monitoring, and the data it provides, is all about context. Systems change, threats change, and so, too, does context. Tuning, therefore, will be ongoing under any continuous monitoring program. This includes, but is not limited to the following:

- Securing vulnerable services
- Ensuring ongoing patch management
- Resolving questionable network traffic issues
- Resolving log analysis issues
- Tuning the monitoring

Securing Vulnerable Services

Even if systems start out secure, they can quickly slip out of configuration. So, continued assessment of systems is critical in finding and remediating those that fall out of configuration. Sometimes tuning must happen at the initial assessment stage. Some services and hosts can be removed, and some can be reconfigured. However, there will usually be some insecure services considered critical by somebody in the organization that will require a decision around remediation measures to take (removal, reconfiguration, isolation of problem service or adding additional protection, such as specialized monitoring).

Ensuring Ongoing Patch Management

All organizations should have a centralized approval process for patch management. Under the program, patches should be approved before being applied. You should run reports to identify what devices have been patched and what software is being used. This means developing automated workflow processes that work with the speed of business. Too often, such efforts are seen as being inhibitors. Also, watch out for small applications that don't have their own centralized patch management solution, making it harder to automate some areas of patch management and workflow.

Resolving Questionable Network Traffic Issues

When you detect traffic that does not match expectations, more tuning may be necessary. If unknown hosts are detected, they probably should be tracked to particular switch ports and cable runs. Many times, the vulnerability assessment identifies the type of host and host name, making the vulnerability easier to find.

A vulnerability scan may show open ports on hosts that are not expected. It is important to try to understand why the port is open before simply closing the port. For example, if port 5800 shows up as being open on a host, it would be best to identify why that port is open. Port 5800 is typically a VNC server. So, ask: Does the organization use VNC for support? If not, do some research to see who installed the software. If there is no apparent explanation for the software installation, scan the rest of the system to see whether it might have been installed maliciously. Also look at the installation directory to see if you can determine when the software was installed.

Resolving Log Analysis Issues

Analyzing log data can reveal a wealth of network tuning opportunities—things that have probably needed to be done for a while that the organization just didn't know about. Don't panic when you see such log data. Do your research and proceed thoughtfully. Moving too quickly on suspicion of a hostile activity could damage evidence, make it difficult to determine the source and make prosecution impossible. If you find that you do have some type of hostile activity going on, you may want to pull in support from other departments or organizations.

Tuning the Monitoring

Certainly not least important is constantly fine-tuning the "continuous monitoring process." Computer systems are constantly changing—and so are the threats. So, the continuous monitoring process must always be changing as well. New technology will come out that will do a better job of automating all of these processes, and the processes will need to be updated. It is an endless cycle.

Reporting

Once you have collected all the information you can about your computing assets and the threats you face (through discovery), analyzed your vulnerabilities and tuned your network by applying patches and hardening services, then you need to document system status, changes and exemptions.

One important report is the number of exemptions or “special consideration” systems. There will always be systems that, for one reason or another, need to be exempted from some control. Keep an eye on these systems and periodically check to see if they still need to get special treatment. Check with the owners of the systems to verify that they are doing their part in manual monitoring. Often system owners are initially diligent at manual monitoring; but, over time, their diligence starts to wane.

Reports should be run on a periodic schedule, and should also be able to be run on an as-needed basis. These reports should be treated as confidential. Some document vulnerabilities, others identify software and hardware. As a group, the reports document what the organization is monitoring. The systems that generate these reports, as well as the reports that contain critical infrastructure information, should be available only to those who need the information and should be secured from everyone else.

9 www.sans.org/critical-security-controls/control.php?id=14

It's a Big Job: Where Do You Start?

Knowing where to start can seem like the most daunting task of all. Remember that this process starts with discovery, which many of your network monitoring tools already provide to some degree.

Before starting any major changes to your program and configurations, run some initial analyses against discovered systems and take time to generate some “before” system state reports. Every organization has stakeholders who want to ensure that money is being well spent, so you’ll want to be able to show system conditions when you started and what progress you have made. Those comparison reports will also help you justify additional resources to continue with the process.

Here are some quick hits to get you going using the tools and brainpower that already exist in your organization:

1. Determine Key Assets.

What computing assets are critical for day-to-day operation, and what assets are most critical from the intellectual property (IP) standpoint? Realistically, every process owner is going to think that his or her process is the most important. In most cases, systems can be put into a few different categories based on criticality. This criticality rating can help identify which systems should get the initial attention when moving into the tuning phase. It will also help with determining how frequently scanning cycles will run.

2. Follow the Quick Wins.

Start with resolving problems that will give you a “Quick Win,” as suggested in the introduction to the 20CSC¹⁰ and then used throughout the document for each of the critical areas. A Quick Win is a solution that will provide real security improvement that can be implemented with a minimum of cost (including implementation time). A Quick Win is generally not a total solution, but it will make real security improvements quickly. For organizations with a large list of action items, the concept of Quick Wins can boost morale and demonstrate progress to management.

Another strategy would be to prioritize the critical controls from the 20CSC paper that require or use a baseline. Those are the controls that can be most easily monitored continuously. Start there, and you can jump start your continuous monitoring efforts.

3. Prioritize Alerts.

At some point in the initial decision-making process, the organization needs to determine time requirements for alerting on different systems and events. For example, an IDS event indicating that a root shell was established on a critical host is much more important than an alert concerning a connection to a closed port on a firewall.

¹⁰ www.sans.org/critical-security-controls/guidelines.php

It's a Big Job: Where Do You Start? (CONTINUED)

We already mentioned the importance of alerts for failed logins and an increase in outbound blocked traffic, which assume that egress filtering is in place. As important as alerts are, it is also important to avoid superfluous alerting. Your system will be subjected to many attacks, and the vast majority of them will not be successful. But, if an IDS detects a successful attack, that is one that should be alerted on. Any attacks on extremely sensitive servers may be worthy of special consideration and possibly of an actual alert. If an organization must run a service that they know is vulnerable, custom-tailored monitoring and alerting of those specific attacks would be a good idea. Each organization has its own special circumstances that will determine what to alert on. In general, alerting should be limited to actual problems. You don't want to turn your alerting system into an IT version of a car alarm that goes off any time somebody walks too close.

Conclusion

We often hear the terms *real time* or *near realtime*, suggesting that if an organization doesn't deal with the problem immediately, the security battle is lost. Sure, catching a perpetrator in seconds would be great; but, if the problem is resolved before damage is done, that's pretty good. Finding out when you read about it in the paper is really bad.

The goal of continuous monitoring starts with prevention: Knowing your network, systems and state—and continuously patching them and managing configurations. This is more than half the battle. Should those defenses fail, continuous monitoring against threats will make that detection time as short as practically possible. Develop some realistic time frames for your organization, using the 20CSC guidelines for initial goals and metrics to test response time.

Continuous monitoring is not a fuzzy, hard-to-reach goal. Most enterprise organizations have many of the primary tools in place. What organizations need are automated processes and tools that can compare what is known about system state and traffic against deviations from what is expected in regular-enough intervals to reduce risk and response time and improve overall compliance and risk posture through continuous improvements. Continuous monitoring is achievable in most enterprise organizations with the technical, managerial and operational experts that already exist.

About the Author

Jerry Shenk currently serves as a senior analyst for the SANS Institute and is senior security analyst for Windstream Communications. He operates out of Ephrata, PA. Since 1984, Jerry has consulted with companies and financial and educational institutions on issues of network design, security, forensic analysis and penetration testing. His experience spans small home-office systems to global networks. Along with some vendor-specific certifications and a CISSP, Jerry holds six GIAC certifications, all completed with honors: GCIA, GCIH, GCFW, GSNA, GPEN and GCFA. Five of his certifications are GOLD certifications.

SANS would like to thank its sponsor:





Upcoming SANS Training

[Click Here for a full list of all Upcoming SANS Events by Location](#)

SANS SEC460: Enterprise Threat Beta	San Diego, CAUS	Jan 08, 2018 - Jan 13, 2018	Live Event
SANS Security East 2018	New Orleans, LAUS	Jan 08, 2018 - Jan 13, 2018	Live Event
SANS Amsterdam January 2018	Amsterdam, NL	Jan 15, 2018 - Jan 20, 2018	Live Event
Northern VA Winter - Reston 2018	Reston, VAUS	Jan 15, 2018 - Jan 20, 2018	Live Event
SEC599: Defeat Advanced Adversaries	San Francisco, CAUS	Jan 15, 2018 - Jan 20, 2018	Live Event
SANS Dubai 2018	Dubai, AE	Jan 27, 2018 - Feb 01, 2018	Live Event
SANS Las Vegas 2018	Las Vegas, NVUS	Jan 28, 2018 - Feb 02, 2018	Live Event
SANS Miami 2018	Miami, FLUS	Jan 29, 2018 - Feb 03, 2018	Live Event
Cyber Threat Intelligence Summit & Training 2018	Bethesda, MDUS	Jan 29, 2018 - Feb 05, 2018	Live Event
SANS Scottsdale 2018	Scottsdale, AZUS	Feb 05, 2018 - Feb 10, 2018	Live Event
SANS London February 2018	London, GB	Feb 05, 2018 - Feb 10, 2018	Live Event
SANS Southern California- Anaheim 2018	Anaheim, CAUS	Feb 12, 2018 - Feb 17, 2018	Live Event
SANS Secure India 2018	Bangalore, IN	Feb 12, 2018 - Feb 17, 2018	Live Event
Cloud Security Summit & Training 2018	San Diego, CAUS	Feb 19, 2018 - Feb 26, 2018	Live Event
SANS Dallas 2018	Dallas, TXUS	Feb 19, 2018 - Feb 24, 2018	Live Event
SANS Brussels February 2018	Brussels, BE	Feb 19, 2018 - Feb 24, 2018	Live Event
SANS Secure Japan 2018	Tokyo, JP	Feb 19, 2018 - Mar 03, 2018	Live Event
SANS New York City Winter 2018	New York, NYUS	Feb 26, 2018 - Mar 03, 2018	Live Event
CyberThreat Summit 2018	London, GB	Feb 27, 2018 - Feb 28, 2018	Live Event
SANS London March 2018	London, GB	Mar 05, 2018 - Mar 10, 2018	Live Event
SANS Cyber Defense Initiative 2017	OnlineDCUS	Dec 12, 2017 - Dec 19, 2017	Live Event
SANS OnDemand	Books & MP3s OnlyUS	Anytime	Self Paced