Sponsored by
Tripwire

Secure Configuration Management Demystified

August 2012

A SANS Whitepaper

Written by: Dave Shackleford

Introduction

Locking down and properly configuring security parameters on servers, workstations and network infrastructure platforms has long been a mainstay of best practices in information security. This typically includes securing local services and applications, setting permissions on files and directories, improving access control parameters such as SSH, limiting privileged user access, and tweaking many other specific settings.

As far back as 2004, Gartner's John Pescatore predicted the need for host-based security to protect an expanding perimeter.[1] Meanwhile, in that same year, a study released by British Telecom and Gartner concluded that up to 65 percent of successful external attacks were directly related to configuration errors.[2] Seven years later, in 2011, Gartner considered secure configuration management as a must-have rather than a nice-to-have control, ranking it No. 1 on the list of server protection priorities.[3] Most recently, in 2012, SANS lists secure configuration for systems, servers and end points as a third critical control, and secure configuration for network and security devices as a tenth critical control.[4,5]

Despite this and other expert guidance, many organizations still lack a sound configuration management strategy to both lock down hosts and network devices and maintain a rigorous security posture over time. According to a recent *InformationWeek* survey of more than 900 IT professionals, enforcing security policies ranked as the second most difficult to achieve, even as they rank vulnerability and patch management among the most beneficial practices.[6] Why? Because we IT professionals think it's too hard! How do we even know what systems we have, let alone what state they're in? Say we do make a gold build that we maintain on our critical systems. What happens when we want to make changes to those systems in a hurry? How do we prevent others from inadvertently changing systems when they administer them? How and where do we automate and integrate to simplify these processes across a mixed enterprise environment? Another question fundamental to secure configuration management is, "How much can we lock down a system before we're impacting system productivity?"

Fortunately, most aspects of configuration management are much more automated today than in the past, but such difficulties persist, resulting in long windows of vulnerability through which attackers infiltrate, embed and spread roots. According to the latest Verizon Business report, misconfigurations and known vulnerabilities have been the top means by which attacks have been successful in the majority of breaches investigated.[7]

Aligning configuration management and patching to the business involves careful measurement and planning and can, even then, still be a complex undertaking when managing exceptions. However, secure configuration management controls need not be as difficult as everyone thinks it is. In this paper, we show how to use secure configuration concepts to reduce the overall attack surface, bring better coordination among groups within IT and elsewhere, and ultimately reduce the risk to your business by continuously improving the IT environment.

---

1 www.gartner.com/DisplayDocument?doc_cd=119940

2 British Telecommunications, "Security and Business Continuity Solutions From BT: Thriving in the Age of the Digital Networked Economy," 2004.

3 Neil MacDonald and Peter Firstbrook, "How To Devise a Server Protection Strategy," December 2011. www.gartner.com/id=1866915

4 www.sans.org/critical-security-controls/control.php?id=3

5 www.sans.org/critical-security-controls/control.php?id=10

6 http://reports.informationweek.com/abstract/21/8815/Security/research-2012-strategic-security-survey.html

7 www.verizonbusiness.com/resources/reports/rp_data-breach-investigations-report-2011_en_xg.pdf

Vulnerabilities and Breaches

Most types of attacks, whether executed through a network channel or application flaw, have one thing in common: They focus on configurable hosts and network devices. The host is where data is stored, applications and services are running, and vulnerabilities are likely to occur that can be exploited to get further into the network. Manipulation of network traffic in transit is possible, but that simply implies that an attacker has some level of access to network conduits, perhaps by gaining access to a switch or through some other means. An attacker who has bypassed network access controls to gain entrance to a network ecosystem cannot hide "on the network." It simply isn't possible. They can, however, hide on hosts.

Ultimately, there must be an endpoint that is compromised to provide shelter to an attacker. In addition, unless Denial of Service (DoS) is the intent, an attacker's goal is always related to a host or network device in some way: data in a database or stored in a file system, user accounts that provide access to hosts, application access, routing configuration, and so on.

Lack of secure configuration management has led to such compromise scenarios repeatedly. Table 1 provides several examples of publicized data breaches and outage scenarios where poor implementation of configuration controls contributed to the compromise or outage.

Breach/Outage	Details
Utah Department of Health (2012)[8]	More than 228,000 Social Security numbers were exposed due to configuration errors.
DropBox (2011)[9]	Configuration errors allowed any DropBox user to access other users' data.
Sony (2011)	Although not confirmed, it is believed that a combination of illicit access by authorized employees and contractors, unpatched and poorly configured systems, and firmware access from PS3 consoles resulted in personal information of 77 million PlayStation network customers being compromised.
Microsoft (2009)[10]	The Bing search engine was offline for a period of time due to configuration errors.
Heartland Payments (2008)	After an initial SQL injection flaw was exploited, intruders leveraged poorly configured systems to hide in the internal network and install malware and custom sniffing software.
TJX (2007)	Poorly configured wireless access points, as well as unsecured and poorly configured in-store kiosks, were used to gain access the internal network.

Table 1: Public Breaches Resulting from Poor Configuration Controls

8 www.cio.in/news/worst-security-errors-2012-284922012

9 http://techcrunch.com/2011/06/20/dropbox-security-bug-made-passwords-optional-for-four-hours/

10 http://mashable.com/2009/12/04/bing-outage-error/

Data from the 2011 and 2012 Verizon Data Breach Investigations Reports (DBIRs) strongly indicates that configuration management and poor control over configuration defaults, upgrades and general system hardening contribute to most data breaches. In the 2011 breach report,[11] several major statistics stood out in relation to poor configuration practices:

- **92% of attacks were not highly difficult.** This implies that attackers were able to easily gain access to systems, applications and data—most likely via existing vulnerabilities and configuration issues. As a professional penetration tester, the author can confirm that many organizations are still missing numerous system security patches and are using default or easily-guessed credentials to protect data.
- **76% of all data was compromised from servers.** This corresponds to the previous item (and configuration concerns in general). Servers are the most prevalent platform in most enterprises, and they require the most overall configuration focus and monitoring.
- **96% of breaches were avoidable through simple or intermediate controls.** Although explicit mention of which controls should have been implemented is omitted from the report, it is implied that stronger patching and configuration of systems would likely have prevented many successful attacks.

The 2011 report also provides a succinct list of areas in which mitigation efforts should be focused, mentioning that organizations should "ensure essential controls are met." Again, no specifics are really defined, but given that the top threat actions are "hacking" and "malware," there's a broad array of controls that could apply, including configuration management and use of antimalware agents, firewalls, network access controls and other such utilities.

Ultimately, according to the Verizon reports, 257 breaches were related to exploitation of default or guessable credentials, 189 were related to malware that tampered with or disabled security controls, and 65 were tied to abuse of system access and/or privileges.

So what changed in 2012? Not much!

- 96% of attacks were not highly difficult.
- 94% of data compromised involved servers.
- 97% of breaches were avoidable through simple or intermediate controls.

Overall, 44 percent of breaches involved exploitation of default or guessable credentials. The most prevalent hacking vectors in the 2012 report were combinations of remote access services, such as Remote Desktop Protocol (RDP) and Virtual Network Computing (VNC), combined with weak and/or default credentials.[12]

As long as organizations treat sound configuration management and change monitoring practices as a set of security controls and processes that is nice to have instead of essential to implement, these types of statistics are likely to continue.

---

11 www.verizonbusiness.com/resources/reports/rp_data-breach-investigations-report-2011_en_xg.pdf

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

What Is Secure Configuration Management?

When considering secure configuration management, the biggest challenge may be simple nomenclature—everyone calls this something different. SANS calls these controls "secure configurations" for servers, workstations, and network and security devices. The National Security Agency (NSA) calls the same thing "patch management" and "baseline management," among other things. The Payment Card Industry Data Security Standards (PCI DSS) compliance standard refers to requirements for secure configurations such as "Do not use vendor-supplied defaults for system passwords and other security parameters" and "Develop and maintain secure systems and applications," without explicitly using the terms "configuration management" or "secure configuration." In NIST 800-53 version 4, currently available as a draft for review, there is a general category of security controls labeled "configuration management."

Fundamentally, all of these terms refer to the same concepts and topics. Throughout this document we call this process *secure configuration management*.

What exactly is secure configuration management? In a nutshell, secure configuration management is the technical application and maintenance of security policy on systems, applications and network devices. There are several distinct disciplines that exist within the realm of secure configuration management, including the following:

- **Configuration Management Planning and Management.** This aspect of configuration management is primarily concerned with developing the configuration management plan, including who will manage it, what type of Configuration Management Database (CMDB) will exist, and what types of tools and processes will be employed.
- **Configuration Identification.** During the Configuration Identification stage, specific Configuration Items (CIs) are defined for individual platforms. These are what will be implemented, measured and maintained over time. For systems and network devices, these controls may be driven by compliance mandates such as the PCI DSS and the Federal Information Security Management Act (FISMA), as well as National Institute of Standards and Technology (NIST) documents (e.g., 800-53) and the Center for Internet Security (CIS) and Defense Information Systems Agency (DISA) guides.
- **Configuration Control.** Once CIs are implemented, organizations can deploy a monitoring system for the controls that focuses primarily on change detection and serves as an holistic change management program. This stage is implemented at two levels. The first application is within the overall CI database and the specific CIs for systems. When a configuration template needs to be changed or updated, it should be accommodated within the change management framework. The second way this phase is managed is at a local level, where changes need to be monitored, implemented, and in some cases, rolled back or remediated when the change is either unplanned or due to malicious activity.

- **Configuration Status Accounting.** This configuration management process focuses on defining and documenting baseline configurations and then maintaining the baselines over time.
- **Configuration Verification and Audit (AKA Configuration Assessment).** Configurations for systems and network devices need to be audited and monitored, with detailed reports prepared for comparison with existing and planned baselines. Performing such audits continuously is best for maintaining adequate security, but regularly scheduled routine audits can also be implemented.
- **Configuration Remediation.** After configuration assessments are performed, there will usually be a variety of configuration items that need to be corrected or changed in some way to meet policy and compliance specifications. IT operations teams usually perform these changes with input from audit and IT security groups. Remediation can be issued in the form of written guidance, as well as through more automated scripts and workflows that are kicked off after change approval.

Challenges with Secure Configuration Management

In an *InformationWeek* survey of CISOs, "enforcing security policies" was ranked as the No. 2 "biggest security challenge," and the biggest overall challenge was managing complexity.[13] It helps to break down the majority of security and compliance-related controls maintained by security and operations teams into the following categories:

- **Physical host operating systems.** Physical servers and workstations make up the bulk of IT infrastructure. The operating platform that runs on these physical systems has a number of distinct controls including access controls, operating system tuning parameters and file-level controls for data storage.
- **Virtual operating platforms and guest operating systems.** As more organizations realize the cost savings and operational benefits of virtualizing systems, an increasing amount of virtual technology will be deployed. In addition to traditional operating system controls, these systems contain their own platform-specific code and virtual network infrastructure that require controls to be applied. To that end, virtualization brings about an entirely new layer of complexity, with a staggering array of new configuration items that should be evaluated and locked down.
- **Applications and databases.** Applications and databases have their own specific sets of controls to maintain and monitor. Web servers have access controls, file and directory permissions, as well as customized database calls and scripts, and databases have controls for data presentation and retention, as well as stored procedures and other database implementation-specific aspects.
- **Network devices and infrastructure.** Network devices such as routers, switches and firewalls all have specific controls that apply to segmentation of network traffic, inspection of traffic and encryption of traffic.

For each of these four categories of infrastructure elements, there are significant numbers of configuration controls that can be applied to implement a specific security or compliance goal. All organizations will have basic physical security controls around hosts, but financial services organizations might have many additional layers of network security controls compared to a small software firm. The software firm, however, might have many more complex code repositories, databases and application-specific controls for some systems. In the case of virtualization, other layers of configuration controls are called for, such as locking down virtual machine files and securing the hypervisor.

The sum of these controls constitutes a security or compliance posture that operates within the framework of internal policies and external regulations and mandates. While that sounds simple enough, managing these controls and processes has been difficult for organizations to follow holistically. For example, they may be managing their virtual hosts, but upgrades to the host platform may affect the virtual guest systems that run on it, leading to modifications in operating system parameters. Over time, the virtual hosts slip out of compliance with the CIS Windows 2008 Server benchmark that serves as their corporate configuration standard—most likely without the organization's knowledge. Because this standard is what the auditing program has been built on, such change could lead to a failed audit or compliance violation as well as risky exposure.

13 http://reports.informationweek.com/abstract/21/8815/Security/research-2012-strategic-security-survey.html

Whether in an IT environment that is managed internally, virtually or in the cloud, the primary culprit behind changes that "get out of hand" is, for most organization, a lack of visibility. This lack of visibility can be seen in a number of forms, including the following:

- **Network visibility.** Different subnets are segregated in such a way that certain systems cannot be identified from other parts of the network, or configurations are only visible to network operations teams.
- **Application visibility.** Ports and services may be unavailable to certain networks, groups or individuals.
- **Configuration visibility.** Certain groups and individuals may not have the access rights to determine a system or application's configuration options.
- **Communication issues.** Certain groups may intentionally or accidentally withhold information from other teams, which leads to a lack of visibility. Without the proper visibility, many teams will not be up-to-date on what systems are deployed, what state those systems are in, and how these changes and discrepancies may affect other parts of the organization. This lack of communication presents the window of opportunity attackers look for to exploit.

Aside from lack of visibility, many organizations have a tendency to leave the definition, management and governance of configuration management to individual groups or silos within IT. Rather than defining a central configuration management team that coordinates tools, processes and policies related to configuration management, many organizations let individual teams of administrators and engineers choose how and when to update and manage system configuration details. Organizations need a coordinated configuration management database or a well-maintained catalog of CIs that all teams refer to for baselines and updates. In most cases, not having such a resource is attributed to staff members not having enough time or senior management lacking commitment. Many IT professionals also feel there is not enough operational capacity to implement and manage a configuration management solution. These silos also lead to gaps in coverage of critical systems that open the doors to exploitation.

There are also challenges with the technical aspects of configuration management. Tools that can both implement and manage configuration details for most platforms tend to come in agent-based formats, with some offering non-agent-based assessment mechanisms as well. Most organizations prefer not to install yet another agent on their systems if possible, due to the perception of reduced performance, more overhead, and possible compatibility issues with other existing software and operating system components.

On the other hand, the agentless options tend to be implemented as remote scanning tools, which may have limited ability to implement configuration changes even when run with administrative credentials. The Department of Homeland Security explains in much more detail the pros and cons of using agents and agentless assessment in its Continuous Asset Evaluation, Situational Awareness, and Risk Scoring Reference Architecture Report (CAESARS).[14]

14 www.dhs.gov/xlibrary/assets/fns-caesars.pdf

From the security perspective, any of these tools should be highly aware of changes to systems, as these are almost always the leading indicators of, or precursors to, security issues such as illicit logins, exploited vulnerabilities and malware installation. Ultimately, however, the only true way to monitor changes to server systems on a continuous basis, as recommended by the SANS 20 Critical Security Controls and other best practices, is by installing an agent. Yet, network configuration management tools lend themselves more readily to remote scanning-based assessments, because many of those operating platforms do not support agents in the first place. For this reason, it's highly likely that a combination of both agented and agentless assessments will be required.

In fact, SANS Critical Control 3 reinforces this point in its current language regarding metrics: "Any of these changes to a computer system must be detected within 24 hours and notification performed by alerting or sending e-mail notification to a list of enterprise administrative personnel. Systems must block installation, prevent execution, or quarantine unauthorized software within one additional hour, alerting or sending e-mail when this action has occurred. Every 24 hours after that point, the system must alert or send e-mail about the status of the system until it has been removed from the network or remediated. While the 24-hour and one-hour timeframes represent the current metric to help organizations improve their state of security, in the future organizations should strive for even more rapid alerting and isolation, with notification about unauthorized changes sent within two minutes."[15]

Along those lines is another challenge related to the *general timeliness* of getting CI information and assessments. A monthly "megascan" is not only invasive, but it is also not sufficiently timely for critical CIs. A newly created illegal service or unexpected telnet session might be recorded by log or SIEM systems, but if not acted upon quickly, it could cause significant damage. Instead, if the security configuration process is real-time and self-reliant (and sometimes self-correcting), it actively serves as a last line of defense.

15 www.sans.org/critical-security-controls/control.php?id=3

Foundational Elements of SANS Controls 3 and 10 and How to Implement Them

In its current iteration, the SANS Critical Security Controls rank secure configuration management controls in two locations among the top 10 for criticality. Why? In a nutshell, the NSA, Australian government, and numerous other contributors determined that these are some of the more likely vectors by which a security breach may occur. In addition, the NSA ranks the mitigation value provided by these controls as "very high" for system hardware and software configuration (Control 3) and "high/medium" for secure configuration on network devices (Control 10).

The following steps provide a general outline for an implementation process and guidelines to achieve secure configuration management in both categories:

Discovery

As mentioned earlier, the first stumbling block to secure configuration management is visibility. This is where a proper discovery process comes in. By leveraging network scanners, system-level scanners and specialized scanning tools that can peruse files and storage infrastructure, physical and virtual assets can be discovered and placed into an inventory. Once this inventory has been developed and validated, a process is needed to continuously discover new assets (or changes in assets) as soon as they are online or shortly thereafter.

Hardening

With a sound inventory in place, organizations need to determine a set of configuration items they want or need to develop and maintain. Most organizations are free to develop their own internal standards that meet policies and compliance guidelines, whereas some others (such as Federal agencies) may be required to adhere to standards like the DISA Secure Technical Implementation Guides (STIGs) or NSA secure configuration guides. Using agent-based and/or agentless technology, organizations will need to apply the configuration standard to systems and then begin assessing the new configuration for changes or deviations from policy. Figure 1 illustrates the various aspects of securely configuring systems.

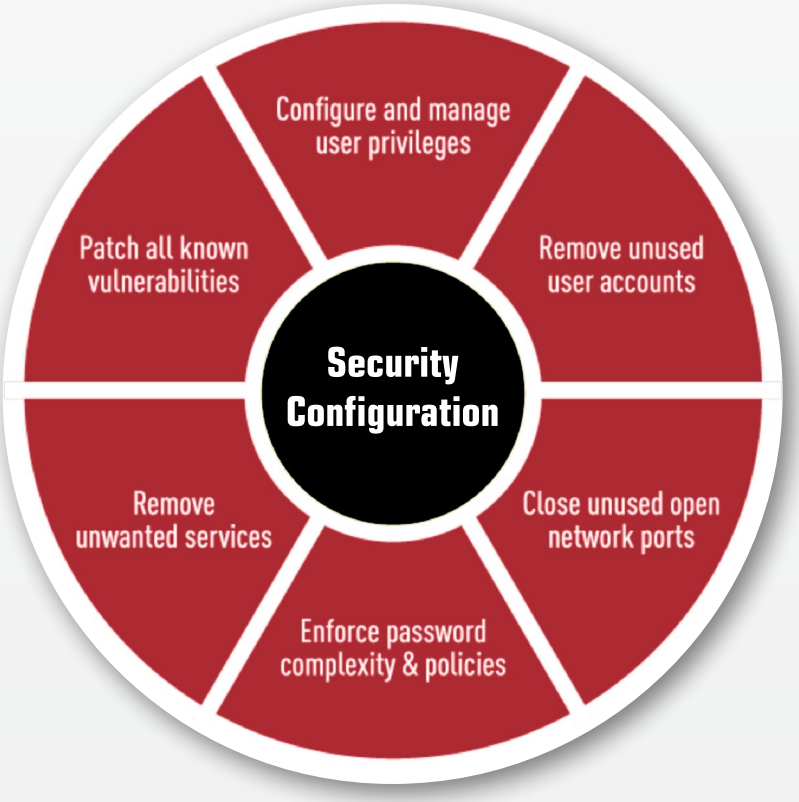

Figure 1: Stages of Security Hardening (adapted with permission from www.sans.org/reading_room/analysts_program/20CriticalControls.pdf)

Management

Once configurations have been applied, systems have a tendency to "shift and drift" from their original postures, which can lead to woefully insecure configurations that are susceptible to attacks. Therefore, ongoing management of the secure configuration tools and CIs is critical. There are several other aspects to proper management of security configuration:

- **Exceptions.** The exception process for secure configuration application needs to be tied closely to change management. When exceptions are required to existing and planned configuration baselines, there should be extensive documentation of why the exception is needed, along with an approval workflow.
- **Waivers.** Waivers should only be issued after the exception process has been successfully followed and documented. In addition, waivers should be revisited regularly to revisit mitigation options.
- **Risk management.** Regularly evaluating the risks of your current configuration templates and CIs should be a part of the secure configuration management lifecycle.
- **Lifecycle.** Developing and applying a configuration to servers, workstations and network devices requires a cyclical approach that ensures configuration controls are updated, patches are applied and systems are evaluated regularly.

Secure configuration management may allow security and operations teams to truly measure configuration baselines and then chart improvements over time, providing a rich source of meaningful metrics. For example, systems may be 50 percent compliant with a chosen configuration standard at the time of initial assessment, and remediation changes can modify this posture to 70 percent over a period of time dictated by project, change cycles and assigned priorities.

Conclusion

Given the prevalence of breaches directly related to or involving configuration failures for systems and network devices, it's important that organizations get back to basics and look at the fundamentals of securing systems and network devices. Implementing a secure configuration management program is essential to properly securing systems and data—and meeting compliance mandates.

Unfortunately, many organizations, especially with the advent of virtualization and private cloud technology, are unable to keep pace with the growing complexity in their environments. The situation only gets worse without rigorous configuration audits and controls. A continuous audit of configuration state and new systems coming online is the best way to manage and monitor all the different infrastructure components that comprise today's complex environments.

Whatever management structure is put in place, organizations must consider changes and updates to their environments. Structured change management processes that do exist must consider that planned and unplanned changes occur out-of-band. Such changes ultimately lead to the majority of most security breaches.

In out-of-band or unmanaged systems, patches aren't applied, configurations are modified, and new virtual systems are brought online without any stewardship from risk managers. Using automation and policy to continuously monitor for changes and auditing configurations are the key means for mitigating risk in the first place, according to the SANS 20 controls and other authority documents. Secure configuration management can also help organizations know when something is occurring on their networks that shouldn't be. Ultimately, proper secure configuration management should lead to continuous overall improvement in risk posture through reduced vulnerability and misconfiguration metrics and cleaner networked systems.

About the Author

Dave Shackleford, founder and principal consultant with Voodoo Security, is a SANS analyst, instructor and course author, as well as a GIAC technical director. He has consulted with hundreds of organizations in the areas of security, regulatory compliance, and network architecture and engineering. He is a VMware *vExpert* and has extensive experience designing and configuring secure virtualized infrastructures. He has previously worked as CSO for Configuresoft and CTO for the Center for Internet Security. Dave is the co-author of *Hands-On Information Security* from Course Technology as well as the author of the upcoming book *Virtualization Security* from Sybex. Recently, Dave co-authored the first published course on virtualization security for the SANS Institute, and leads its virtualization and cloud security curriculum. Dave currently serves on the board of directors at the Technology Association of Georgia's Information Security Society and the SANS Technology Institute.

SANS would like to thank its sponsor: