Twenty Critical Controls for Effective Cyber Defense: Consensus Audit Guidelines

Version 2.3: November 13, 2009

Update: Added NIST SP 800-53 Revision 3 mapping to each control, and updated appendix to include each area of direct mapping between 20 Critical Controls and 800-53 Rev 3 Priority 1 controls. Also, added metrics and tests for each of the automatable controls (the first 15). Finally, added an appendix summarizing the attack types that motivated the development of each control.

INTRODUCTION

Securing our nation against cyber attacks has become one of the nation's highest priorities. To achieve this objective, networks, systems, and the operations teams that support them must vigorously defend against a variety of threats, both internal and external. Furthermore, for those attacks that are successful, defenses must be capable of detecting, thwarting, and responding to follow-on attacks on internal enterprise networks as attackers spread inside a compromised network.

A central tenet of the US Comprehensive National Cybersecurity Initiative (CNCI) is that "offense must inform defense." In other words, knowledge of actual attacks that have compromised systems provides the essential foundation on which to construct effective defenses. The US Senate Homeland Security and Government Affairs Committee moved to make this same tenet central to the Federal Information Security Management Act in drafting the U.S. ICE Act of 2009 (the new FISMA). That new proposed legislation calls upon federal agencies to (and on the White House to ensure that they):

> *"monitor, detect, analyze, protect, report, and respond against known vulnerabilities, attacks, and exploitations" and "continuously test and evaluate information security controls and techniques to ensure that they are effectively implemented."*

Because federal agencies do not have unlimited money, current and past federal CIOs and CISOs have agreed that the only rational way they can hope to meet these requirements is to jointly establish a prioritized baseline of information security measures and controls that can be continuously monitored through automated mechanisms. In addition, most agencies have highly interconnected systems and information requiring an enterprise approach to cyber security since cyber attacks exploit the weak areas in an enterprise to gain access to other enterprise capabilities. As we look to the future, it is also clear that ongoing initiatives within the federal government will continue to expand interconnectivity across agencies to better support citizens and internal government operations. Security vulnerabilities in one area of a particular federal agency can become the path to compromise other parts of the federal enterprise. It is essential

that a prioritized set of security controls be established that can be applied across agency enterprise environments and potentially across the federal government.

This consensus document of 20 crucial controls is designed to begin the process of establishing that *prioritized baseline of information security measures and controls* that can be applied across federal enterprise environments. The consensus effort that has produced this document has identified 20 specific technical security controls that are viewed as effective in blocking currently known high-priority attacks, as well as those attack types expected in the near future. Fifteen of these controls can be monitored, at least in part, automatically and continuously. The consensus effort has also identified a second set of five controls that are essential but that do not appear to be able to be monitored continuously or automatically with current technology and practices. Each of the 20 control areas includes multiple individual subcontrols, each specifying actions an organization can take to help improve its defenses.

The control areas and individual subcontrols described focus on various technical aspects of information security, with a primary goal of supporting organizations in prioritizing their efforts in defending against today's most common and damaging computer and network attacks. Outside of the technical realm, a comprehensive security program should also take into account numerous additional areas of security, including overall policy, organizational structure, personnel issues (e.g., background checks, etc.), and physical security. To help maintain focus, the controls in this document do not deal with these important, but non-technical, aspects of information security. Organizations should build a comprehensive approach in these other aspects of security as well, but overall policy, organization, personnel, and physical security are outside of the scope of this document.

In summary, the guiding principles used in devising these control areas and their associated subcontrols include:

- Defenses should focus on addressing the most common and damaging attack activities occurring today, and those anticipated in the near future.
- Enterprise environments must ensure consistent controls across an enterprise to effectively negate attacks
- Defenses should be automated where possible, and periodically or continuously measured using automated measurement techniques where feasible.
- To address current attacks occurring on a frequent basis against numerous organizations, a variety of specific technical activities should be undertaken to produce a more consistent defense.

Additionally, the controls are designed to support agencies and organizations that currently have different levels of information security capabilities. To help organizations focus on achieving a sound baseline of security and then improve beyond that baseline, certain subcontrols have been categorized as follows:

- *Quick Wins:* These fundamental aspects of information security can help an organization rapidly improve its security stance generally without major procedural, architectural, or

technical changes to its environment. It should be noted, however, that a *Quick Win* does not necessarily mean that these subcontrols provide comprehensive protection against the most critical attacks. The intent of identifying *Quick Win* areas is to highlight where security can be improved rapidly. These items are identified in this document with the label of "QW."

- *Improved Visibility and Attribution:* These subcontrols focus on improving the process, architecture, and technical capabilities of organizations so that organizations can monitor their networks and computer systems, gaining better visibility into the IT operations. Attribution is associated with determining which computer systems, and potentially which users, are generating specific events. Such improved visibility and attribution support organizations in detecting attack attempts, locating the points of entry for successful attacks, identifying already-compromised machines, interrupting infiltrated attackers' activities, and gaining information about the sources of an attack. In other words, these controls help to increase an organization's situational awareness of its environment. These items are labeled as "Vis/Attrib."
- *Hardened Configuration and Improved Information Security Hygiene:* These aspects of various controls are designed to improve the information security stance of an organization by reducing the number and magnitude of potential security vulnerabilities as well as improving the operations of networked computer systems. This type of control focuses on protecting against poor security practices by system administrators and end users that could give an adversary an advantage in attacking target systems. Control guidelines in this category are formulated with the understanding that a well managed network is typically a much harder target for computer attackers to exploit. Throughout this document, these items are labeled as "Config/Hygiene."
- *Advanced:* These items are designed to further improve the security of an organization beyond the other three categories. Organizations already following all of the other controls should focus on this category. Items in this category are simply called "Advanced."

In general, organizations should examine all 20 control areas against their current status and develop an agency-specific plan to implement the controls as a critical component of an overall security program. Ultimately, organizations should strive to implement each control area, applying all of the subcontrols within each control, working from QW, through Vis/Attrib and Config/Hygiene, up to Advanced. However, as a start, organizations with limited information security programs may want to address the "Quick Wins" aspects of the controls in order to make rapid progress and to build momentum within their information security program.

Many of these controls can be implemented and measured using existing tools found in many enterprises. Other controls can be fulfilled using commercial or, in some cases, free, open-source software. Still others may require an investment in new enterprise tools and personnel expertise.

Each control area also includes a metric section that provides detailed information about the specific timing and related objectives associated with the most important elements of the given control area. Furthermore, each control area also includes a test section that provides information about how organizations can evaluate their implementation of each control metric. These tests were devised to support automation wherever possible, so that organizations could achieve reliable, scalable, and continuous measurements of their adherence to the controls and related metrics.

Why This Project Is So Important: Gaining Agreement among CISOs, CIOs and IGs, with Technical Requirements for System Administrators and Security Personnel

Federal Chief Information Security Officers (CISOs) and Chief Information Officers (CIOs) are charged with improving the state of information security across the federal government. Moreover, they are spending increasing amounts of money to secure their systems. However, the complexity of securing their systems is enormous, and therefore there is a need to focus attention and resources on the most critical risk (and therefore the highest payoff) areas. In addition, CISOs and CIOs want and need specific guidance that can be consistently applied across their agencies enterprise-wide and upon which their performance in improving security can be consistently and fairly evaluated. At the same time, Federal Inspectors General (IGs) and auditors want to ensure that CIOs and CISOs are doing what is necessary to secure systems, but IGs and auditors, too, need specific guidance on how to measure security. And, finally, technical personnel associated with information security operations and system administration require a specific set of technical activities that will aid them in defending against current and near-term attack vectors.

This document is a first step toward providing specific guidelines that CISOs, CIOs, IGs, and various Computer Emergency Response Teams can adopt and provide to their technical system administration and information security personnel to ensure their agency systems have the most critical baseline security controls in place. The controls take advantage of the knowledge gained in analyzing the myriad attacks that are being actively and successfully launched against federal systems and our nation's industrial base systems.

This effort also takes advantage of the success and insights from the development and usage of standardized concepts for identifying, communicating, and documenting security-relevant characteristics/data. These standards areas include the common identification of vulnerabilities, definition of secure configurations, inventory of systems and platforms, vulnerability severity, and identification of application weaknesses. These standards have emerged over the last decade through collaborative research and deliberation between government, academia, and industry. While still evolving, these efforts have made their way into commercial solutions and government, industry, and academic usage. Perhaps most visible of these has been the Security Content Automation Program (SCAP), sponsored by NIST, which was mandated for the Federal Desktop Core Configuration (FDCC). SCAP utilizes mature standardizations to clearly define common security nomenclature and evaluation criteria for

vulnerability, patch, and configuration measurement and is intended for adoption by automated tools. It is recommended that automated tools used to implement or verify security controls identified in this document employ SCAP or similar standardization efforts for clearly defined nomenclature and evaluation criteria not covered by SCAP.

Relationship of the 20 Critical Controls to NIST Guidelines

The National Institute of Standards and Technology (NIST) has produced excellent security guidelines that provide a very comprehensive set of security controls in NIST Special Publication 800-53, revision 3. This document by contrast seeks to identify a subset of security control activities that CISOs, CIOs and IGs can focus on as their top, shared priority for cyber security based on attacks occurring today and those anticipated in the near future. As noted above, the 20 Critical Controls only address principally technical control areas. However, they also address many critical operational controls as identified in NIST Special Publication 800-53. It is recommended that 800-53 be used by agencies to ensure that they have assessed and implemented an appropriate set of management controls. Each control described in this document includes a mapping to the specific corresponding 800-53 area(s) where the two documents are consistent in their requirements. This mapping, which is included within each individual control and also appears in totality as an appendix to this document, demonstrates that the 20 Critical Controls are a proper subset of 800-53 Priority 1 items.

The authors of this document recommended that agency CIOs and CISOs assess the 20 Critical Controls as a baseline set of "Common Controls" for their agency as defined by 800-53. The basis for this recommendation is the fact that the consensus process used to develop the 20 Critical Controls as well as pilot efforts by the State Department have validated that the controls correlate to the highest technical and operational threat areas for federal agency enterprise environment (as well as private sector enterprise environments). Within the guidance of 800-53, the 20 Critical Controls can be viewed as necessary to address agency "high water mark" when assessing the enterprise-wide potential security risk of confidentiality, integrity or availability of interconnected systems and information within the agency's enterprise environment. Once this agreement is reached by the CIO and CISO, it is recommended that the 20 Critical Controls would be the foundation for technical and operational controls within an agency. Similarly, the 20 Critical Controls would also serve as a primary basis for future security audits and evaluations. If an agency CIO and CISO determined that their environment warranted additional controls, the processes provided in 800-53 should be used to identify additional required controls. Based on the overwhelming consensus from security experts who contributed to this document, it is viewed as unlikely that an agency with Internet connectivity would determine that the 20 Critical Controls were not applicable to their agency. However, if an agency CIO and CISO determined that some of the 20 Critical Controls were not applicable to an agency, it is recommended that this also be documented using processes outlined in 800-53.

Document Contributors

What makes this document effective is that it reflects knowledge of actual attacks and defines controls that would have stopped those attacks from being successful. To construct the document, the following types of people have provided first-hand knowledge and input regarding how computer and network attacks are being carried out and the defensive techniques that are most important in thwarting attacks:

1. Blue team members inside the Department of Defense who are often called in when military commanders find their systems have been compromised and who perform initial incident response services on impacted systems
2. Blue team members who provide services for non-DoD government agencies who identify prior intrusions while conducting vulnerability assessment activities
3. US-CERT and other non-military incident response employees and consultants who are called upon by civilian agencies and companies to identify the most likely method by which systems and networks were compromised
4. Military investigators who fight cyber crime
5. The FBI and other police organizations that investigate cyber crime
6. Cybersecurity experts at US Department of Energy laboratories and federally funded research and development centers
7. DoD and private forensics experts who analyze computers that have been infected
8. Red team members in DoD tasked with finding ways of circumventing military cyber defenses during their exercises
9. Civilian penetration testers who test civilian government and commercial systems to determine how they can be penetrated with the goal of better understanding risk and implementing better defenses
10. Federal CIOs and CISOs who have intimate knowledge of cyber attacks

Additionally, input from over one hundred other collaborators has been incorporated into the current version of the document. To assemble these Top 20 Critical Controls, these contributors first identified the most prevalent and damaging attack types and scenarios, so that appropriate defenses could be identified. These attacks are described in the introduction to each individual control in a section titled “How do attackers exploit the lack of this control?” Furthermore, Appendix B of this document provides a list of each of the attack types that fueled the development of the Top 20 Critical Controls.

The Twenty Critical Controls

These 20 critical security controls were agreed upon by knowledgeable individuals from the groups listed above. The list includes 15 controls that can be validated at least in part in an automated manner and five that must be validated manually. It is important to note that the 20 control categories are *not* presented in order of priority. The process of gathering these specific controls and subcontrols focused on identifying the highest priority defenses and

represent a subset of controls found in other audit guidelines and documents. Each of the 20 categories is important and offers high-priority techniques for thwarting real-world attacks.

Critical Controls Subject to Automated Collection, Measurement, and Validation:

1. Inventory of Authorized and Unauthorized Devices
2. Inventory of Authorized and Unauthorized Software
3. Secure Configurations for Hardware and Software on Laptops, Workstations, and Servers
4. Secure Configurations for Network Devices such as Firewalls, Routers, and Switches
5. Boundary Defense
6. Maintenance, Monitoring, and Analysis of Security Audit Logs
7. Application Software Security
8. Controlled Use of Administrative Privileges
9. Controlled Access Based on Need to Know
10. Continuous Vulnerability Assessment and Remediation
11. Account Monitoring and Control
12. Malware Defenses
13. Limitation and Control of Network Ports, Protocols, and Services
14. Wireless Device Control
15. Data Loss Prevention

Additional Critical Controls (not directly supported by automated measurement and validation):

16. Secure Network Engineering
17. Penetration Tests and Red Team Exercises
18. Incident Response Capability
19. Data Recovery Capability
20. Security Skills Assessment and Appropriate Training to Fill Gaps

In the pages that follow, each of these controls is described more fully. Descriptions include how attackers currently exploit the lack of the control, detailed subcontrols that describe what an organization needs to do in each area and requirements for measuring these activities, and suggestions regarding how standardized measurements can be applied. As pilot implementations are completed and agencies gain experience with automation, it is expected that the document will be expanded into a detailed audit guide that agency CIOs can use to ensure they are doing the right things for effective cyber defense and that IGs can use to verify the CIOs' tests.

Insider Threats vs. Outsider Threats

A quick review of the critical controls may lead some readers to think that they are heavily focused on outsider threats and may, therefore, not fully deal with insider attacks. In reality, the insider threat is well covered in these controls in two ways. First, specific controls such as

maintenance of security audit logs, control of administrative privileges, controlled access based on need to know, data loss prevention, and effective incident response all directly address the key ways that insider threats can be mitigated. Second, the insider and outsider threats sometimes merge as outsiders penetrate security perimeters and effectively become "insiders." All of the controls that limit unauthorized access within the organization work to mitigate both insider and outsider threats. It is important to note that these controls are meant to deal with multiple kinds of computer attackers, including but not limited to malicious internal employees and contractors, independent individual external actors, organized crime groups, terrorists, and nation state actors, as well as mixes of these different threats. While these controls are designed to provide protection against each of these threats, very sophisticated, well-funded actors such as nation states may sometimes employ attack techniques that require extreme defenses which go beyond the scope of this document.

These controls are not limited to blocking only the initial compromise of systems, but also address detecting already-compromised machines, and preventing or disrupting attacker's actions. The defenses identified through these controls deal with decreasing the initial attack surface by hardening security, identifying already-compromised machines to address long-term threats inside an organization's network, controlling super-user privileges on systems, and disrupting attackers' command-and-control of implanted malicious code. Figure 1 illustrates the scope of different kinds of attacker activities that these controls are designed to help thwart.

The rings of Figure 1 represent the actions computer attackers often take against target machines. These actions include initially compromising a machine to establish a foothold by exploiting one or more vulnerabilities. Attackers can then maintain long-term access on a system, often by creating accounts, subverting existing accounts, or altering the software on the machine to include backdoors and rootkits. Attackers with access to machines can also cause damage, which could include stealing, altering, or destroying information; impairing the system's functionality to jeopardize its business effectiveness or mission; or using it as a jump-off point for compromise of other systems in the environment. Where these rings overlap, attackers have even more ability to compromise sensitive information or cause damage.

The various defensive strategies located outside of each set of rings in the figure are covered throughout the controls described in this document. Defenses in any of the rings helps to limit the abilities of attackers, but improved defenses are required across all three rings and their intersections. It is important to note that the Twenty Critical Controls for Effective Cyber Defense are designed to help improve defenses across each of these rings, rather than to merely prevent initial compromise.

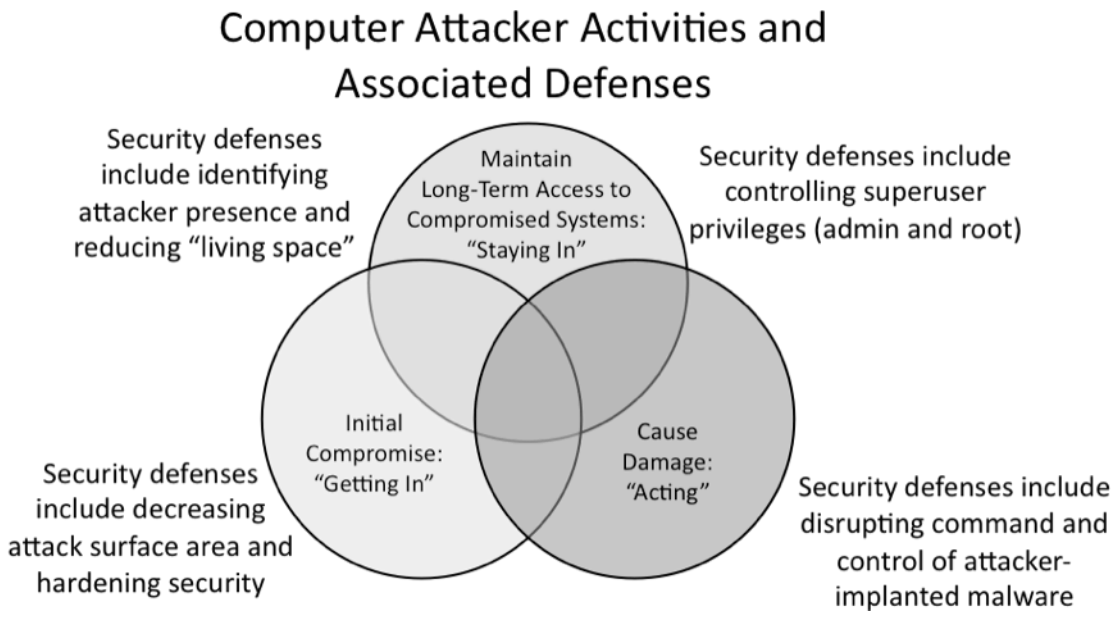

Figure 1: Types of Computer Attacker Activities These Controls Are Designed to Help Thwart

Relationship to Other Federal Guidelines, Recommendations, and Requirements

These controls are meant to reinforce and prioritize some of the most important elements of the guidelines, standards, and requirements put forth in other US Government documentation, such as NIST Special Publication 800-53: *Recommended Security Controls for Federal Information Systems*, SCAP, FDCC, FISMA, and Department of Homeland Security Software Assurance documents. These guidelines do not conflict with such recommendations. In fact, the guidelines set forth are a proper subset of the recommendations of NIST SP 800-53, designed so that organizations can focus on a specific set of actions associated with current threats and computer attacks they face every day. A draft of the mapping of individual controls in this document to specific recommendations of NIST SP 800-53 is included in Appendix A.

Periodic and Continual Testing of Controls

Each control included in this document describes a series of tests that organizations can conduct on a periodic or, in some cases, continual basis to ensure that appropriate defenses are in place. One of the goals of the tests described is to provide as much automation of testing as possible. By leveraging standardization efforts and repositories of content like SCAP, these automated test suites and scripts can be highly sharable between organizations, consistent to a large extent, and easily used by auditors for validation. A key element to support automation of measurement is the management infrastructure of the enterprise network. Well managed networks tend to have enterprise tools for remotely gathering, analyzing, and updating the configuration of workstations, servers, and network equipment on a fine-grained basis.

It is important to note that, at various phases of the tests described in the controls, human testers are needed to set up tests or evaluate results in a fashion that cannot be automated. The testers responsible for measuring such controls must be trusted individuals, as the test may require them to access sensitive systems or data in the course of their tests. Without appropriate authorization, background checks, and possibly clearance, such tests may be impossible. Such tests should also be supervised or reviewed by appropriate agency officials well versed in lawful monitoring and analysis of Information Technology systems as well as regulatory requirements for protecting sensitive Personally Identifiable Information.

Future Evolution of the Twenty Critical Controls for Effective Cyber Defense

The consensus effort to define critical security controls is an evolving effort. In fact, changing technology and changing attack patterns will necessitate future changes even after the current set of controls has been finalized. In a sense, this will be a living document moving forward, but the controls described in this version are a solid start on the quest to make fundamental computer security defenses a well understood, repeatable, measurable, scalable, and reliable process throughout the federal government.

Critical Control 1: Inventory of Authorized and Unauthorized Devices

How do attackers exploit the lack of this control?

Many criminal groups and nation states deploy systems that continuously scan address spaces of target organizations waiting for new, unprotected systems to be attached to the network. The attackers also look for laptops not up to date with patches because they are not frequently connected to the network. One common attack takes advantage of new hardware that is installed on the network one evening and not configured and patched with appropriate security updates until the following day. Attackers from anywhere in the world may quickly find and exploit such systems that are Internet-accessible. Furthermore, even for internal network systems, attackers who have already gained internal access may hunt for and compromise additional improperly secured internal computer systems. Some attackers use the local nighttime window to install backdoors on the systems before they are hardened.

Additionally, attackers frequently look for experimental or test systems that are briefly connected to the network but not included in the standard asset inventory of an organization. Such experimental systems tend not to have as thorough security hardening or defensive measures as other systems on the network. Although these test systems do not typically hold sensitive data, they offer an attacker an avenue into the organization, and a launching point for deeper penetration.

How can this control be implemented, automated, and its effectiveness measured?

An accurate and up-to-date inventory, controlled by active monitoring and configuration management, can reduce the chance of attackers finding unauthorized and unprotected systems to exploit.

1. QW: Deploy an automated asset inventory discovery tool and use it to build a preliminary asset inventory of systems connected to the enterprise network. Both active tools that scan through network address ranges, and passive tools that identify hosts based on analyzing their traffic should be employed.
2. Vis/Attrib: Maintain an asset inventory of all systems connected to the network and the network devices themselves, recording at least the network addresses, machine name(s), purpose of each system, an asset owner responsible for each device, and the department associated with each device. The inventory should include every system that has an IP address on the network, including, but not limited to desktops, laptops, servers, network equipment (routers, switches, firewalls, etc.), printers, Storage Area Networks, Voice-over-IP telephones, etc.
3. Vis/Attrib: Ensure that network inventory monitoring tools are operational and continuously monitoring, keeping the asset inventory up to date on a real-time basis,

looking for deviations from the expected inventory of assets on the network, and alerting security and/or operations personnel when deviations are discovered.

4. Config/Hygiene: Secure the asset inventory database and related systems, ensuring that they are included in periodic vulnerability scans and that asset information is encrypted. Limit access to these systems to authorized personnel only, and carefully log all such access. For additional security, a secure copy of the asset inventory may be kept in an off-line system air-gapped from the production network.
5. Config/Hygiene: In addition to an inventory of hardware, organizations should develop an inventory of information assets, which identifies their critical information, and maps critical information to the hardware assets (including servers, workstations, and laptops) on which it is located. A department and individual responsible for each information asset should be identified, recorded, and tracked.
6. Advanced: The organization's asset inventory should include removable media devices, including USB tokens, external hard drives, and other related information storage devices.

Associated NIST SP 800-53 Rev 3 Priority 1 Controls:

CM-8 (a, c, d, 2, 3, 4), PM-5, PM-6

Procedures and tools for implementing and automating this control:

Organizations must first establish information owners and asset owners, deciding and documenting which organizations and individuals are responsible for each component of information and device. Some organizations maintain asset inventories using specific large-scale enterprise commercial products dedicated to the task or they use free solutions to track and then sweep the network periodically for new assets connected to the network. In particular, when effective organizations acquire new systems, they record the owner and features of each new asset, including its network interface MAC address, a unique identifier hard-coded into most network interface cards and devices. This mapping of asset attributes and owner-to-MAC address can be stored in a free or commercial database management system.

Then, with the asset inventory assembled, many organizations use tools to pull information from network assets such as switches and routers regarding the machines connected to the network. Using securely authenticated and encrypted network management protocols, tools can retrieve MAC addresses and other information from network devices that can be reconciled with the organization's asset inventory of servers, workstations, laptops, and other devices.

Going further, effective organizations configure free or commercial network scanning tools to perform network sweeps on a regular basis, such as every 12 hours, sending a variety of different packet types to identify devices connected to the network. Before such scanning can take place, organizations should verify that they have adequate bandwidth for such periodic

scans by consulting load history and capacities for their networks. In conducting inventory scans, scanning tools could send traditional ping packets (ICMP Echo Request), looking for ping responses to identify a system at a given IP address. Because some systems block inbound ping packets, in addition to traditional pings, scanners can also identify devices on the network using TCP SYN or ACK packets. Once they have identified IP addresses of devices on the network, some scanners provide robust fingerprinting features to determine the operating system type of the discovered machine.

In addition to active scanning tools that sweep the network, other asset identification tools passively listen on network interfaces looking for devices to announce their presence by sending traffic. Such passive tools can be connected to switch span ports at critical places in the network to view all data flowing through such switches, maximizing the chance of identifying systems communicating through those switches.

Wireless devices (and wired laptops) may periodically join a network and then disappear making the inventory of currently available systems churn significantly. Likewise, virtual machines can be difficult to track in asset inventories when they are shut down or paused, because they are merely files in some host machine's file system. Additionally, remote machines accessing the network using VPN technology may appear on the network for a time, and then be disconnected from it. Each machine, whether physical or virtual, directly connected to the network or attached via VPN, currently running or shut down, should be included in an organization's asset inventory.

Control 1 Metric:

The system must be capable of identifying any new unauthorized devices that are connected to the network within 24 hours, and alerting or sending email notification to a list of enterprise administrative personnel. The system must automatically isolate the unauthorized system from the network within one hour of the initial alert and send a follow-up alert or e-mail notification when isolation is achieved. 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. The asset inventory database and alerting system must be able to identify the location, department, and other details of where authorized and unauthorized devices are plugged into the network. While the 24 hour and one hour timeframes represent the current metric to help organizations improve their current state of security, in the future, organizations should strive for even more rapid alerting and isolation, with notification about an unauthorized asset connected to the network being sent within two minutes and isolation within five minutes.

Control 1 Test:

To evaluate the implementation of Control 1 on a periodic basis, the evaluation team will connect hardened test systems to at least ten locations on the network, including a selection of subnets associated with DMZs, workstations, and servers. Two of the systems must be included in the asset inventory database, while the other systems are not. The evaluation team must then verify that the systems generate an alert or e-mail notice regarding the newly connected systems within 24 hours of the test machines being connected to the network. The evaluation team must verify that the system provides details of the location of all the test machines connected to the network. For those test machines included in the asset inventory, the team must also verify that the system provides information about the asset owner.

The evaluation team must then verify that the test systems are automatically isolated from the production network within one hour of initial notification and that an e-mail or alert indicating the isolation has occurred. The team must then verify that the connected test systems are isolated from production systems by attempting to ping and use other protocols to access systems on the production network and checking that connectivity is not allowed.

Critical Control 2: Inventory of Authorized and Unauthorized Software

How do attackers exploit the lack of this control?

Computer attackers deploy systems that continuously scan address spaces of target organizations looking for vulnerable versions of software that can be remotely exploited. Some attackers also distribute hostile web pages, document files, media files, and other content via their own web pages or otherwise trustworthy third-party sites. When unsuspecting victims access this content with a vulnerable browser or other client-side program, attackers compromise their machines, often installing backdoor programs and bots that give the attacker long-term control of the system. Some sophisticated attackers may use zero-day exploits, which take advantage of previously unknown vulnerabilities for which no patch has yet been released by the software vendor. Without proper knowledge or control of the software deployed in an organization, defenders cannot properly secure their assets.

Without the ability to inventory and control which programs are installed and allowed to run on their machines, enterprises make their systems more vulnerable. Such poorly controlled machines are more likely to be either running software that is unneeded for business purposes, introducing potential security flaws, or running malware introduced by a computer attacker after system compromise. Once a single machine has been exploited, attackers often use it as a staging point for collecting sensitive information from the compromised system and from other systems connected to it. In addition, compromised machines are used as a launching point for movement throughout the network and partnering networks. In this way, attackers may

quickly turn one compromised machine into many. Organizations that do not have complete software inventories are unable to find systems running vulnerable or malicious software to mitigate problems or root out attackers.

How can this control be implemented, automated, and its effectiveness measured?

1. QW: Devise a list of authorized software that is required in the enterprise for each type of system, including servers, workstations, and laptops of various kinds and uses.
2. Vis/Attrib: Deploy software inventory tools throughout the organization covering each of the operating system types in use, including servers, workstations, and laptops. The software inventory system should track the version of the underlying operating system as well as the applications installed on it. Furthermore, the tool should record not only the type of software installed on each system, but also its version number and patch level. The tool should also monitor for unauthorized software installed on each machine. This unauthorized software also includes legitimate system administration software installed on inappropriate systems where there is no business need for it.
3. Advanced: Deploy software white-listing technology that allows systems to run only approved applications and prevents execution of all other software on the system.

Associated NIST SP 800-53 Rev 3 Priority 1 Controls:

CM-1, CM-2 (2, 4, 5), CM-3, CM-5 (2, 7), CM-7 (1, 2), CM-8 (1, 2, 3, 4, 6), CM-9, PM-6, SA-6, SA-7

Procedures and tools for implementing and automating this control:

Commercial software and asset inventory tools are widely available and in use in many enterprises today. The best of these tools provide an inventory check of hundreds of common applications used in enterprises, pulling information about the patch level of each installed program to ensure that it is the latest version and leveraging standardized application names, such as those found in the Common Platform Enumeration (CPE) specification.

Features that implement whitelists and blacklists of programs allowed to run or blocked from executing are included in many modern end-point security suites. Moreover, commercial solutions are increasingly bundling together anti-virus, anti-spyware, personal firewall, and host-based Intrusion Detection Systems and Intrusion Prevention Systems (IDS and IPS), along with software white listing and black listing. In particular, most endpoint security solutions can look at the name, file system location, and/or cryptographic hash of a given executable to determine whether the application should be allowed to run on the protected machine. The most effective of these tools offer custom whitelists and blacklists based on executable path, hash, or regular expression matching. Some even include a graylist function that allows administrators to define rules for execution of specific programs only by certain users and at certain times of day, and blacklists based on specific signatures.

Control 2 Metric:

The system must be capable of identifying unauthorized software, by detecting either an attempt to install it or execute it, notifying enterprise administrative personnel within 24 hours through an alert or email. 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. While the 24 hour and one hour timeframes represent the current metric to help organizations improve their current state of security, in the future, organizations should strive for even more rapid alerting and isolation, with notification about unauthorized software being sent within two minutes and isolation within five minutes.

Control 2 Test:

To evaluate the implementation of Control 2 on a periodic basis, the evaluation team must move a benign software test program that is not included in the authorized software list to ten systems on the network. Two of the systems must be included in the asset inventory database, while the other systems are not. The evaluation team must then verify that the systems generate an alert or e-mail notice regarding the new software within 24 hours. The team must also verify that the alert or e-mail is received within one additional hour indicating that the software has been blocked or quarantined. The evaluation team must verify that the system provides details of the location of each machine with this new test software, including information about the asset owner.

The evaluation team must then verify that the software is blocked by attempting to execute it, and verifying that the software is not allowed to run.

Critical Control 3: Secure Configurations for Hardware and Software on Laptops, Workstations, and Servers

How do attackers exploit the lack of this control?

On both the Internet and internal networks that attackers have already compromised, automated computer attack programs constantly search target networks looking for systems that were configured with vulnerable software installed the way that it was delivered from manufacturers and resellers, thereby being immediately vulnerable to exploitation. Default configurations are often geared to ease-of-deployment and ease-of-use and not security,

leaving some systems exploitable in their default state. Attackers attempt to exploit both network-accessible services and browsing client software using such techniques.

Defenses against these automated exploits include procuring computer and network components with the secure configurations already implemented, deploying such pre-configured hardened systems, updating these configurations on a regular basis, and tracking them in a configuration management system.

How can this control be implemented, automated, and its effectiveness measured?

1. QW: System images must have documented security settings that are tested before deployment, approved by an agency change control board, and registered with a central image library for the agency or multiple agencies. These images should be validated and refreshed on a regular basis (such as every six months) to update their security configuration in light of recent vulnerabilities and attack vectors.
2. QW: Standardized images should represent hardened versions of the underlying operating system and the applications installed on the system, such as those released by NIST, NSA, DISA, the Center for Internet Security (CIS), and others. This hardening would typically include removal of unnecessary accounts, as well as the disabling or removal of unnecessary services. Such hardening also involves, among other measures, applying patches, closing open and unused network ports, implementing intrusion detection systems and/or intrusion prevention systems, and host-based firewalls.
3. QW: Any deviations from the standard build or updates to the standard build should be documented and approved in a change management system.
4. QW: Government agencies should negotiate contracts to buy systems configured securely out of the box using standardized images, which should be devised to avoid extraneous software that would increase their attack surface and susceptibility to vulnerabilities.
5. QW: The master images themselves must be stored on securely configured servers, with integrity checking tools and change management to ensure only authorized changes to the images are possible. Alternatively, these master images can be stored in off-line machines, air-gapped from the production network, with images copied via secure media to move them between the image storage servers and the production network.
6. Config/Hygiene: At least once per month, run assessment programs on a varying sample of systems to measure the number that are and are not configured according to the secure configuration guidelines.
7. Config/Hygiene: Utilize file integrity checking tools on at least a weekly basis to ensure that critical system files (including sensitive system and application executables, libraries, and configurations) have not been altered. All alterations to such files should be automatically reported to security personnel. The reporting system should have the ability to account for routine and expected changes, highlighting unusual or unexpected alterations.

8. Config/Hygiene: Implement and test an automated configuration monitoring system that measures all secure configuration elements that can be measured through remote testing, using features such as those included with SCAP-compliant tools to gather configuration vulnerability information. These automated tests should analyze both hardware and software changes, network configuration changes, and any other modifications affecting security of the system.
9. Config/Hygiene: Provide senior executives with charts showing the number of systems that match configuration guidelines versus those that do not match, illustrating the change of such numbers month by month for each organizational unit.

Associated NIST SP 800-53 Rev 3 Priority 1 Controls:

CM-1, CM-2 (1, 2), CM-3 (b, c, d, e, 2, 3), CM-5 (2), CM-6 (1, 2, 4), CM-7 (1), SA-1 (a), SA-4 (5), SI-7 (3), PM-6

Procedures and tools for implementing this control:

Organizations can implement this control by developing a series of images and secure storage servers for hosting these standard images. Then, commercial and/or free configuration management tools can be employed to measure the settings of managed machines' operating system and applications to look for deviations from the standard image configurations used by the organization. Some configuration management tools require that an agent be installed on each managed system, while others remotely login to each managed machine using administrator credentials. Either approach or combinations of the two approaches can provide the information needed for this control.

Control 3 Metric:

The system must be capable of identifying any changes to an official hardened image that may include modifications to key files, services, ports, configuration files or any software installed on the system. Modifications include deletion, changes or additions of new software to any part of the operating systems, services or applications running on the system. The configuration of each system must be checked against the official master image database to verify any changes to secure configurations that would impact security. Any of these changes to a computer system must be detected within 24 hours and notification performed by alerting or sending email 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 current state of security, in the future,

organizations should strive for even more rapid alerting and isolation, with notification about unauthorized changes being sent within two minutes and installation and execution blocked within five minutes.

Control 3 Test:

To evaluate the implementation of Control 3 on a periodic basis, an evaluation team must move a benign test system that does not contain the official hardened image, containing additional services, ports and configuration files changes, onto the network. This must be performed on ten different random segments, using either real or virtual systems. The evaluation team must then verify that the systems generate an alert or e-mail notice regarding the changes to the software within 24 hours. It is important that the evaluation team verify that all unauthorized changes have been detected. The team must also verify that the alert or e-mail is received within one additional hour indicating that the software has been blocked or quarantined. The evaluation team must verify that the system provides details of the location of each machine with the unauthorized changes, including information about the asset owner.

The evaluation team must then verify that the software is blocked by attempting to execute it and verifying that it is not allowed to run. In addition to these tests, two additional tests must be performed:

1) File integrity checking tools must be run on a regular basis. Any changes to critical operating system, services and configuration files must be checked on an hourly basis. Any changes must be blocked and follow the above email notification process.
2) System scanning tools that check for open ports, services, software version, patch levels and configuration files must be run on a daily basis. Any changes must be blocked and follow the above email notification process.

Critical Control 4: Secure Configurations for Network Devices such as Firewalls, Routers, and Switches

How do attackers exploit the lack of this control?

Attackers take advantage of the fact that network devices may become less securely configured over time as users demand exceptions for specific and temporary business needs, the exceptions are deployed, and those exceptions are not undone when the business need is no longer applicable. Making matters worse, in some cases, the security risk of the exception is never properly analyzed, nor is this risk measured against the associated business need. Attackers search for electronic holes in firewalls, routers, and switches and use those to

penetrate defenses. Attackers have exploited flaws in these network devices to gain access to target networks, redirect traffic on a network (to a malicious system masquerading as a trusted system), and to intercept and alter information while in transmission. Through such actions, the attacker gains access to sensitive data, alters important information, or even uses one compromised machine to pose as another trusted system on the network.

How can this control be implemented, automated, and its effectiveness measured?

1. QW: Compare firewall, router, and switch configuration against standard secure configurations defined for each type of network device in use in the organization. The security configuration of such devices should be documented, reviewed, and approved by an agency change control board. Any deviations from the standard configuration or updates to the standard configuration should be documented and approved in a change control system.
2. QW: At network interconnection points, such as Internet gateways, inter-agency connections, and internal network segments with different security controls, implement ingress and egress filtering to allow only those ports and protocols with a documented business need. All other ports and protocols besides those with an explicit need should be blocked with default-deny rules by firewalls, network-based IPSs, and/or routers.
3. QW: Network devices that filter unneeded services or block attacks (including firewalls, network-based Intrusion Prevention Systems, routers with access control lists, etc.) should be tested under laboratory conditions with each given organization's configuration to ensure that these devices exhibit failure behavior in a closed/blocking fashion under significant loads with traffic including a mixture of legitimate, allowed traffic for that configuration intermixed with attacks at line speeds.
4. Config/Hygiene: All new configuration rules beyond a baseline-hardened configuration that allow traffic to flow through network security devices, such as firewalls and network-based IPSs, should be documented and recorded in a configuration management system, with a specific business reason for each change, a specific individual's name responsible for that business need, and an expected duration of the need. At least once per quarter, these rules should be reviewed to determine whether they are still required from a business perspective. Expired rules should be removed.
5. Config/Hygiene: Network filtering technologies employed between networks with different security levels (firewalls, network-based IPS tools, and routers with ACLs) should be deployed with capabilities to filter IPv6 traffic. Even if IPv6 is not explicitly used on the network, many operating systems today ship with IPv6 support activated, and therefore filtering technologies need to take it into account.
6. Config/Hygiene: Network devices should be managed using two-factor authentication and encrypted sessions. Only true two-factor authentication mechanisms should be used, such as a password and a hardware token, or a password and biometric device. Requiring two different passwords for accessing a system is not two-factor authentication.

7. Advanced: The network infrastructure should be managed across network connections that are separated from the business use of that network, relying on separate VLANs or preferably relying on entirely different physical connectivity for management sessions for network devices.

Associated NIST SP 800-53 Rev 3 Priority 1 Controls:

AC-4 (7, 10, 11, 16), CM-1, CM-2 (1), CM-3 (2), CM-5 (1, 2, 5), CM-6 (4), CM-7 (1, 3), IA-2 (1, 6), IA-5, IA-8, RA-5, SC-7 (2, 4, 5, 6, 8, 11, 13, 14, 18), SC-9

Procedures and tools for implementing this control:

Some organizations use commercial tools that evaluate the rule set of network filtering devices to determine whether they are consistent or in conflict, providing an automated sanity check of network filters and search for errors in rule sets or ACLs that may allow unintended services through the device. Such tools should be run each time significant changes are made to firewall rule sets, router ACLs, or other filtering technologies.

Control 4 Metric:

The system must be capable of identifying any changes to network devices including routers, switches, firewalls, IDS and IPS systems. These changes include any modifications to key files, services, ports, configuration files or any software installed on the device. Modifications include deletions, changes or additions of new software to any part of the device configuration. The configuration of each system must be checked against the official master image database to verify any changes to secure configurations that would impact security. This includes both operating system and configuration files. Any of these changes to a device must be detected within 24 hours and notification performed by alerting or sending email notification to a list of enterprise personnel. If possible, devices must prevent changes to the system and send an e-mail indicating the change was not successful. Every 24 hours after that point, the system must alert or send e-mail about the status of the system until it is investigated and/or remediated.

Control 4 Test:

To evaluate the implementation of Control 4 on a periodic basis, an evaluation team must make a change to each type of network device plugged into the network. At a minimum, routers, switches, and firewalls need to be tested. If they exist, IPS, IDS, and other network devices must be included. Backups must be made prior to making any changes to critical network devices. It is critical that changes do not impact or weaken the security of the device. Acceptable changes include but are not limited to making a comment or adding a duplicate entry in the

configuration. The change must be performed twice for each critical device. The evaluation team must then verify that the systems generate an alert or e-mail notice regarding the changes to the device within 24 hours. It is important that the evaluation team verify that all unauthorized changes have been detected and have resulted in an alert or e-mail notification. The evaluation team must verify that the system provides details of the location of each device, including information about the asset owner. While the 24 hour timeframe represents the current metric to help organizations improve their current state of security, in the future, organizations should strive for even more rapid alerting and isolation, with notification about unauthorized configuration changes in network devices being sent within two minutes.

If appropriate an additional test must be performed on a daily basis to ensure that other protocols such as IPv6 are properly being filtered.

Critical Control 5: Boundary Defense

How do attackers exploit the lack of this control?

Attackers focus on exploiting systems that they can reach across the Internet, which include not only DMZ systems, but also workstation and laptop computers that pull content from the Internet through network boundaries. Threats such as organized crime groups and nation states use configuration and architectural weaknesses found on perimeter systems, network devices, and Internet-accessing client machines to gain initial access into an organization. Then, with a base of operations on these machines, attackers often pivot to get deeper inside the boundary to steal or change information or to set up a persistent presence for later attacks against internal hosts. Additionally, many attacks occur between business partner networks, sometimes referred to as extranets, as attackers hop from one organization's network to another, exploiting vulnerable systems on extranet perimeters.

To control the flow of traffic through network borders and to police its content looking for attacks and evidence of compromised machines, boundary defenses should be multi-layered, relying on firewalls, proxies, DMZ perimeter networks, and network-based Intrusion Prevention Systems and Intrusion Detection Systems.

It should be noted that boundary lines between internal and external networks are diminishing through increased interconnectivity within and between organizations, as well as the rapid rise in deployment of wireless technologies. These blurring lines sometimes allow attackers to gain access inside networks while bypassing boundary systems. However, even with this blurring, effective security deployments still rely on carefully configured boundary defenses that separate networks with different threat levels, different sets of users, and different levels of control. Even with the blurring of internal and external networks, effective multi-layered

defenses of perimeter networks help to lower the number of successful attacks, allowing security personnel to focus on attackers who have devised methods to bypass boundary restrictions.

How can this control be implemented, automated, and its effectiveness measured?

The boundary defenses included in this control build on Critical Control 4, with these additional recommendations focused on improving the overall architecture and implementation of both Internet and internal network boundary points. Internal network segmentation is central to this control because once inside a network, many intruders attempt to target the most sensitive machines. Usually, internal network protections are not set up to defend against an internal attacker. Setting up even a basic level of security segmentation across the network and protecting each segment with a proxy and a firewall will greatly reduce the intruders' access to the other parts of the network.

1. QW: Organizations should deny communications with (or limit data flow to) known malicious IP addresses (blacklists) or limit access to trusted sites (whitelists). Periodically, test packets from bogon source IP addresses should be sent into the network to verify that they are not transmitted through network perimeters. Lists of bogon addresses (unroutable or otherwise unused IP addresses) are publicly available on the Internet from various sources, and indicate a series of IP addresses that should not be used for legitimate traffic traversing the Internet.
2. QW: Deploy IDS sensors on Internet and extranet DMZ systems and networks that look for unusual attack mechanisms and detect compromise of these systems. These IDS sensors may detect attacks through the use of signatures, network behavior analysis, or other mechanisms to analyze traffic.
3. QW: On DMZ networks, monitoring systems (which may be built-in to the IDS sensors or deployed as a separate technology) should be configured to record at least packet header information, and preferably full packet header and payloads of the traffic destined for or passing through the network border.
4. Vis/Attrib: Define a network architecture that clearly separates internal systems from DMZ systems and extranet systems. DMZ systems are machines that need to communicate with the internal network as well as the Internet, while extranet systems are systems whose primary communication is with other systems at a business partner.
5. Vis/Attrib: Design and implement network perimeters so that all outgoing web, FTP, and secure shell traffic to the Internet must pass through at least one proxy on a DMZ network. The proxy should support logging individual TCP sessions; blocking specific URLs, domain names, and IP addresses to implement a blacklist; and applying whitelists of allowed sites that can be accessed through the proxy while blocking all other sites.
6. Vis/Attrib: Require all remote login access (including VPN, dial-up, and other forms of access that allow login to internal systems) to use two-factor authentication.

7. Config/Hygiene: All devices remotely logging into the internal network should be managed by the enterprise, with remote control of their configuration, installed software, and patch levels.
8. Config/Hygiene: Organizations should periodically scan for back-channel connections to the Internet that bypass the DMZ, including unauthorized VPN connections and dual-homed hosts connected to the enterprise network and to other networks via wireless, dial-up modems, or other mechanisms.
9. Config/Hygiene: To limit access by an insider or malware spreading on an internal network, organizations should devise internal network segmentation schemes to limit traffic to only those services needed for business use across the internal network.
10. Config/Hygiene: Organizations should develop plans for rapidly deploying filters on internal networks to help stop the spread of malware or an intruder.
11. Advanced: Organizations should force outbound traffic to the Internet through an authenticated proxy server on the enterprise perimeter.
12. Advanced: To help identify covert channels exfiltrating data through a firewall, built-in firewall session tracking mechanisms included in many commercial firewalls should be configured to identify long-term TCP sessions that last an unusually long time for the given organization and firewall device, alerting personnel about the source and destination addresses associated with these long-term sessions.

Associated NIST SP 800-53 Rev 3 Priority 1 Controls:

AC-17 (1), AC-20, CA-3, IA-2 (1, 2), IA-8, RA-5, SC-7 (1, 2, 3, 8, 10, 11, 14), SC-18, SI-4 (c, 1, 4, 5, 11), PM-7

Procedures and tools for implementing this control:

One element of this control can be implemented using free or commercial IDSs and sniffers to look for attacks from external sources directed at DMZ and internal systems, as well as attacks originating from internal systems against the DMZ or Internet. Security personnel should regularly test these sensors by launching vulnerability-scanning tools against them to verify that the scanner traffic triggers an appropriate alert. The captured packets of the IDS sensors should be reviewed using an automated script each day to ensure that log volumes are within expected parameters and that the logs are formatted properly and have not been corrupted.

Additionally, packet sniffers should be deployed on DMZs to look for HTTP traffic that bypasses HTTP proxies. By sampling traffic regularly, such as over a 3-hour period once per week, information security personnel search for HTTP traffic that is neither sourced by nor destined for a DMZ proxy, implying that the requirement for proxy use is being bypassed.

To identify back-channel connections that bypass approved DMZs, network security personnel can establish an Internet-accessible system to use as a receiver for testing outbound access. This system is configured with a free or commercial packet sniffer. Then, security personnel

connect a sending test system to various points on the organization's internal network, sending easily identifiable traffic to the sniffing receiver on the Internet. These packets can be generated using free or commercial tools with a payload that contains a custom file used for the test. When the packets arrive at the receiver system, the source address of the packets should be verified against acceptable DMZ addresses allowed for the organization. If source addresses are discovered that are not included in legitimate, registered DMZs, more detail can be gathered by using a traceroute tool to determine the path packets take from the sender to the receiver system.

Control 5 Metric:

The system must be capable of identifying any unauthorized packets sent into a trusted zone or out of a trusted zone and ensure that the packets are properly blocked and/or alerted on. Any unauthorized packets must be detected 24 hours, with the system generating an alert or e-mail for enterprise administrative personnel. Alerts must be sent every hour thereafter until the boundary device is reconfigured. While the 24 hour timeframe represents the current metric to help organizations improve their current state of security, in the future, organizations should strive for even more rapid alerting, with notification about unauthorized packets in a trusted zone being sent within two minutes.

Control 5 Test:

To evaluate the implementation of Control 5 on a periodic basis, an evaluation team must test boundary devices by sending packets from outside any trusted network to ensure that only authorized packets are allowed through the boundary. All other packets must be dropped. In addition, unauthorized packets must be sent from a trusted network to an untrusted network to make sure egress filtering is functioning properly. The evaluation team must then verify that the systems generate an alert or e-mail notice regarding the unauthorized packets within 24 hours. It is important that the evaluation team verify that all unauthorized packets have been detected. The evaluation team must also verify that the alert or e-mail is received within one hour indicating that the unauthorized traffic is now being blocked. The evaluation team must verify that the system provides details of the location of each machine with this new test software, including information about the asset owner. It is also important that the evaluation team test to ensure that the device fails in a in a state where it does not forward traffic when it crashes or becomes flooded.

Critical Control 6: Maintenance, Monitoring, and Analysis of Audit Logs

How do attackers exploit the lack of this control?

Deficiencies in security logging and analysis allow attackers to hide their location, malicious software used for remote control, and activities on victim machines. Even if the victims know that their systems were compromised, without protected and complete logging records, the victim is blind to the details of the attack and to the subsequent actions taken by the attackers. Without solid audit logs, an attack may go unnoticed indefinitely and the particular damages done may be irreversible.

Sometimes logging records are the only evidence of a successful attack. Many organizations keep audit records for compliance purposes but attackers rely on the fact that such organizations rarely look at the audit logs so they do not know that their systems have been compromised. Because of poor or non-existent log analysis processes, attackers sometimes control victim machines for months or years without anyone in the target organization knowing, even though the evidence of the attack has been recorded in unexamined log files.

How can this control be implemented, automated, and its effectiveness measured?

1. QW: Validate audit log settings for each hardware device and the software installed on it, ensuring that logs include a date, timestamp, source addresses, destination addresses, and various other useful elements of each packet and/or transaction. Systems should record logs in a standardized format such as syslog entries or those outlined by the Common Event Expression (CEE) initiative. If systems cannot generate logs in a standardized format, deploy log normalization tools to convert logs into a standardized format.
2. QW: Ensure that all systems that store logs have adequate storage space for the logs generated on a regular basis, so that log files will not fill up between log rotation intervals.
3. QW: System administrators and security personnel should devise profiles of common events from given systems, so that they can tune detection to focus on unusual activity, avoid false positives, more rapidly identify anomalies, and prevent overwhelming analysts with insignificant alerts.
4. QW: All remote access to an internal network, whether through VPN, dial-up, or other mechanism, should be logged verbosely.
5. QW: Operating systems should be configured to log access control events associated with a user attempting to access a resource (e.g., a file or directory) without the appropriate permissions.
6. QW: Security personnel and/or system administrators should run bi-weekly reports that identify anomalies in logs. They should then actively review the anomalies, documenting their findings.

7. Vis/Attrib: Each agency network should include at least two synchronized time sources, from which all servers and network equipment retrieve time information on a regular basis, so that timestamps in logs are consistent.
8. Vis/Attrib: Network boundary devices, including firewalls, network-based IPSs, and inbound and outbound proxies should be configured to log verbosely all traffic (both allowed and blocked) arriving at the device.
9. Vis/Attrib: For all servers, organizations should ensure logs are written to write-only devices or to dedicated logging servers running on separate machines from hosts generating the event logs, lowering the chance that an attacker can manipulate logs stored locally on compromised machines.
10. Advanced: Organizations should deploy a Security Event/Information Management (SEIM) system tool for log aggregation and consolidation from multiple machines and for log correlation and analysis. Deploy and monitor standard government scripts for analysis of the logs, as well as using customized local scripts. Furthermore, event logs should be correlated with information from vulnerability scans to fulfill two goals. First, personnel should verify that the activity of the regular vulnerability scanning tools themselves is logged. And, secondly, personnel should be able to correlate attack detection events with earlier vulnerability scanning results to determine whether the given exploit was used against a known-vulnerable target.

Associated NIST SP 800-53 Rev 3 Priority 1 Controls:

AC-17 (1), AC-19, AU-2 (4), AU-3 (1,2), AU-4, AU-5, AU-6 (a, 1, 5), AU-8, AU-9 (1, 2), AU-12 (2), SI-4 (8)

Procedures and tools for implementing this control:

Most free and commercial operating systems, network services, and firewall technologies offer logging capabilities. Such logging should be activated, with logs sent to centralized logging servers. Firewalls, proxies, and remote access systems (VPN, dial-up, etc.) should all be configured for verbose logging, storing all the information available for logging should a follow-up investigation be required. Furthermore, operating systems, especially those of servers, should be configured to create access control logs when a user attempts to access resources without the appropriate privileges. To evaluate whether such logging is in place, an organization should periodically scan through its logs and compare them with the asset inventory assembled as part of Critical Control 1, to ensure that each managed item actively connected to the network is periodically generating logs.

Analytical programs for reviewing logs can be useful, but the capabilities employed to analyze audit logs is quite wide-ranging, including just a cursory examination by a human. Actual correlation tools can make audit logs far more useful for subsequent manual inspection by people. Such tools can be quite helpful in identifying subtle attacks. However, these tools are neither a panacea nor a replacement for skilled information security personnel and system

administrators. Even with automated log analysis tools, human expertise and intuition are often required to identify and understand attacks.

Control 6 Metric:

The system must be capable of logging all events across the network. The logging must be validated across both network and host-based systems. Any event must generate a log entry that includes a date, timestamp, source address, destination address and other details about the packet. Any activity performed on the network must be logged immediately to all devices along the critical path. When a device detects that it is not capable of generating logs (due to a log server crash or other issue), it must generate an alert or e-mail for enterprise administrative personnel within 24 hours. While the 24 hour timeframe represents the current metric to help organizations improve their current state of security, in the future, organizations should strive for even more rapid alerting, with notification about a logging failure being sent within two minutes.

Control 6 Test:

To evaluate the implementation of Control 6 on a periodic basis, an evaluation team must review the security logs of various network devices, servers, and hosts. At a minimum the following devices must be tested: two routers, two firewalls, two switches, ten servers, and ten client systems. The testing team should use traffic-generating tools to send packets through the systems under analysis to verify that the traffic is logged. This analysis is done by creating controlled, benign events and determining if the information is properly recorded in the logs with key information including a date, timestamp, source address, destination address, and other details about the packet. The evaluation team must verify that the system generates audit logs and, if not, an alert or e-mail notice regarding the failed logging must be sent within 24 hours. It is important that the team verify that all activity has been detected. The evaluation team must verify that the system provides details of the location of each machine, including information about the asset owner.

Critical Control 7: Application Software Security

How do attackers exploit the lack of this control?

Attacks against vulnerabilities in web-based and other application software have been a top priority for criminal organizations in recent years. Application software that does not properly

check the size of user input, fails to sanitize user input by filtering out unneeded but potentially malicious character sequences, or does not initialize and clear variables properly could be vulnerable to remote compromise. Attackers can inject specific exploits, including buffer overflows, SQL injection attacks, and cross-site scripting code to gain control over vulnerable machines. In one attack in 2008, more than 1 million web servers were exploited and turned into infection engines for visitors to those sites using SQL injection. During that attack, trusted websites from state governments and other organizations compromised by attackers were used to infect hundreds of thousands of browsers that accessed those websites. Many more web and non-web application vulnerabilities are discovered on a regular basis.

To avoid such attacks, both internally developed and third-party application software must be carefully tested to find security flaws. For third-party application software, enterprises should verify that vendors have conducted detailed security testing of their products. For in-house developed applications, enterprises must conduct such testing themselves or engage an outside firm to conduct such testing.

How can this control be implemented, automated, and its effectiveness measured?

1. QW: Organizations should protect web applications by deploying web application firewalls that inspect all traffic flowing to the web application for common web application attacks, including but not limited to Cross-Site Scripting, SQL injection, command injection, and directory traversal attacks. For applications that are not web based, deploy specific application firewalls if such tools are available for the given application type.
2. Config/Hygiene: Organizations should test in-house developed and third-party procured web and other application software for coding errors and malware insertion, including backdoors prior to deployment using automated static code analysis software. If source code is not available, these organizations should test compiled code using static binary analysis tools. In particular, input validation and output encoding routines of application software should be carefully reviewed and tested.
3. Config/Hygiene: Organizations should test in-house developed and third-party procured web applications for common security weaknesses using automated remote web application scanners prior to deployment, whenever updates are made to the application, and on a regular recurring basis, such as weekly.
4. Config/Hygiene: For applications that rely on a database, organizations should conduct a configuration review of both the operating system housing the database and the database software itself, checking settings to ensure that the database system has been hardened using standard hardening templates.
5. Config/Hygiene: Organizations should verify that security considerations are taken into account throughout the requirements, design, implementation, testing, and other phases of the application development life cycle of all applications.
6. Config/Hygiene: Organizations should ensure that all software development personnel receive training in writing secure code for their specific development environment.

7. Config/Hygiene: Require that all in-house developed software include white-list filtering capabilities for all data input and output associated with the system. These whitelists should be configured to allow in or out only the types of data needed for the system, blocking other forms of data that are not required.

Associated NIST SP 800-53 Rev 3 Priority 1 Controls:

CM-7, RA-5 (a, 1), SA-3, SA-4 (3), SA-8, SI-3, SI-10

Procedures and tools for implementing this control:

Source code testing tools, web application security scanning tools, and object code testing tools have proven useful in securing application software, along with manual application security penetration testing by testers who have extensive programming knowledge as well as application penetration testing expertise. The Common Weakness Enumeration (CWE) initiative is utilized by many such tools to identify the weaknesses that they find. Organizations can also use CWE to determine which types of weaknesses they are most interested in addressing and removing. A broad community effort to identify the "Top 25 Most Dangerous Programming Errors" is also available as a minimum set of important issues to investigate and address during the application development process. When evaluating the effectiveness of testing for these weaknesses, the Common Attack Pattern Enumeration and Classification (CAPEC) can be used to organize and record the breadth of the testing for the CWEs as well as a way for testers to think like attackers in their development of test cases.

Control 7 Metric:

The system must be capable of detecting and blocking an application-level software attack attempt, and must generate an alert or send e-mail to enterprise administrative personnel within 24 hours of detection and blocking.

All Internet-accessible web applications must be scanned on a weekly or daily basis, alerting or sending e-mail to administrative personnel within 24 hours of completing a scan. If a scan cannot be completed successfully, the system must alert or send e-mail to administrative personnel within one hour indicating that the scan has not completed successfully. Every 24 hours after that point, the system must alert or send e-mail about the status of uncompleted scans, until normal scanning resumes.

Additionally, all high-risk vulnerabilities in Internet-accessible web applications identified by web application vulnerability scanners, static analysis tools, and automated database configuration review tools must be mitigated (by either fixing the flaw or through implementing a compensating control) within fifteen days of discovery of the flaw.

While the 24 hour and one hour timeframes represent the current metric to help organizations improve their current state of security, in the future, organizations should strive for even more rapid alerting, with notification about an application attack attempt being sent within two minutes.

Control 7 Test:

To evaluate the implementation of Control 7 on a monthly basis, an evaluation team must use a web application vulnerability scanner to tests for each type of flaw identified in the regularly updated list of the "25 Most Dangerous Programming Errors". The scanner must be configured to assess all of the organization's Internet-accessible web applications to identify such errors. The evaluation team must verify that the scan is detected within 24 hours and an alert is generated.

In addition to the web application vulnerability scanner, the evaluation team must also run static code analysis tools and database configuration review tools against Internet-accessible applications to identify security flaws on a monthly basis.

The evaluation team must verify that all high-risk vulnerabilities identified by the automated vulnerability scanning tools or static code analysis tools have been remediated or addressed through a compensating control (such as a web application firewall) within 15 days of discovery.

The evaluation team must verify that application vulnerability scanning tools have successfully completed their regular scans for the previous 30 cycles of scanning by reviewing archived alerts and reports to ensure that the scan was completed. If a scan was not completed successfully, the system must alert or send e-mail to enterprise administrative personnel indicating that the scan did not complete successfully. If a scan could not be completed in that timeframe, the evaluation team must verify that an alert or e-mail was generated indicating that the scan did not complete.

Critical Control 8: Controlled Use of Administrative Privileges

How do attackers exploit the lack of this control?

According to some Blue Team personnel as well as investigators of large-scale Personally Identifiable Information (PII) breaches, the misuse of administrator privileges is the number one method for attackers to spread inside a target enterprise. Two very common attacker techniques take advantage of uncontrolled administrative privileges. In the first, a workstation

user is fooled into opening a malicious email attachment, downloading and opening a file from a malicious web site, or simply surfing to a website hosting attacker content that can automatically exploit browsers. The file or exploit contains executable code that runs on the victim's machine either automatically or by tricking the user into executing the attacker's content. If the victim user's account has administrative privileges, the attacker can take over the victim's machine completely and install keystroke loggers, sniffers, and remote control software to find administrator passwords and other sensitive data.

The second common technique used by attackers is elevation of privileges by guessing or cracking a password for an administrative user to gain access to a target machine. If administrative privileges are loosely and widely distributed, the attacker has a much easier time gaining full control of systems, because there are many more accounts that can act as avenues for the attacker to compromise administrative privileges. One of the most common of these attacks involves the domain administration privileges in large Windows environments, giving the attacker significant control over large numbers of machines and access to the data they contain.

How can this control be implemented, automated, and its effectiveness measured?

1. QW: Organizations should inventory all administrative passwords and validate that each person with administrative privileges on desktops, laptops, and servers is authorized by a senior executive and that his/her administrative password has at least 12 semi-random characters, consistent with the Federal Desktop Core Configuration (FDCC) standard.
2. QW: Before deploying any new devices in a networked environment, organizations should change all default passwords for applications, operating systems, routers, firewalls, wireless access points, and other systems to a difficult-to-guess value.
3. QW: Organizations should configure all administrative-level accounts to require regular password changes on a 30-, 60-, or 90-day interval.
4. QW: Organizations should ensure all service accounts have long and difficult-to-guess passwords that are changed on a periodic basis as is done for traditional user and administrator passwords.
5. QW: Passwords for all systems should be stored in a hashed or encrypted format. Furthermore, files containing these encrypted or hashed passwords required for systems to authenticate users should be readable only with superuser privileges.
6. QW: Organizations should ensure that administrator accounts are used only for system administration activities, and not for reading e-mail, composing documents, or surfing the Internet.
7. QW: Through policy and user awareness, organizations should require that administrators establish unique, different passwords for their administrator accounts and their non-administrative accounts. On systems with unsalted passwords, such as Windows machines, this approach can be verified in a password audit by comparing the password hashes of each account used by a single person.

8. QW: Organizations should configure operating systems so that passwords cannot be re-used within a certain time frame, such as six months.
9. Vis/Attrib: Organizations should implement focused auditing on the use of administrative privileged functions and monitor for anomalous behavior (e.g., system reconfigurations during night shift).
10. Vis/Attrib: Organizations should configure systems to issue a log entry and alert when an account is added to or removed from a domain administrators group.
11. Config/Hygiene: All administrative access, including domain administrative access, should utilize two-factor authentication.
12. Config/Hygiene: Remote access directly to a machine should be blocked for administrator-level accounts. Instead, administrators should be required to access a system remotely using a fully logged and non-administrative account. Then, once logged in to the machine without admin privileges, the administrator should then transition to administrative privileges using tools such as sudo on Linux/UNIX, runas on Windows, and other similar facilities for other types of systems.
13. Config/Hygiene: Organizations should conduct targeted spear-phishing tests against both administrative personnel and non-administrative users to measure the quality of their defense against social engineering.
14. Advanced: Organizations should segregate administrator accounts based on defined roles within the organization. For example, "Workstation admin" accounts should only be allowed administrative access of workstations, laptops, etc.

Associated NIST SP 800-53 Rev 3 Priority 1 Controls:

AC-6 (2, 5), AC-17 (3), AC-19, AU-2 (4)

Procedures and tools for implementing this control:

Built-in operating system features can extract lists of accounts with superuser privileges, both locally on individual systems and on overall domain controllers. To verify that users with high-privileged accounts do not use such accounts for day-to-day web surfing and e-mail reading, security personnel could periodically gather a list of running processes in an attempt to determine whether any browsers or e-mail readers are running with high privileges. Such information gathering can be scripted, with short shell scripts searching for a dozen or more different browsers, e-mail readers, and document editing programs running with high privileges on machines. Some legitimate system administration activity may require the execution of such programs over the short term, but long-term or frequent use of such programs with administrative privileges could indicate that an administrator is not adhering to this control.

Additionally, to prevent administrators from accessing the web using their administrator accounts, administrative accounts can be configured to use a web proxy of 127.0.0.1 in some operating systems that allow user-level configuration of web proxy settings. Furthermore, in

some environments, administrator accounts do not require the ability to receive e-mail. These accounts can be created without an e-mail box on the system.

To enforce the requirement for password length of 12 or more characters, built-in operating system features for minimum password length can be configured, which prevent users from choosing short passwords. To enforce password complexity (requiring passwords to be a string of pseudo-random characters), built-in operating system settings or third-party password complexity enforcement tools can be applied.

Control 8 Metric:

The system must be configured to comply with password policies at least as stringent as those described in the controls above. Additionally, security personnel must be notified via an alert or e-mail within 24 hours of the addition of an account to a super user group, such as a domain administrator. Every 24 hours after that point, the system must alert or send e-mail about the status of administrative privileges until the unauthorized change has been corrected or authorized through a change management process. While the 24 hour timeframes represent the current metric to help organizations improve their current state of security, in the future, organizations should strive for even more rapid alerting, with notification about new additions to super-use groups being sent within two minutes.

Control 8 Test:

To evaluate the implementation of Control 8 on a periodic basis, an evaluation team must verify that the organization's password policy is enforced by creating a temporary, disabled, limited privilege test account on ten different systems and then attempting to change the password on the account to a value that does not meet the organization's password policy. The selection of these systems must be as random as possible and include a cross-section of the organization's systems and locations. After completion of the test, this account must be removed. Furthermore, the evaluation team must add a temporary disabled test account to a super user group (such as a domain administrator group) to verify that an alert or e-mail is generated within 24 hours. After this test, the account must be removed from the group and disabled.

Finally, on a periodic basis, the evaluation team must run a script that determines which browser and e-mail client programs are running on a sample of ten test systems, including five clients and five servers. Any browsers or mail client software running with Windows administrator or Linux/Unix UID 0 privileges must be identified.

Critical Control 9: Controlled Access Based on Need to Know

How do attackers exploit the lack of this control?

Some organizations do not carefully identify and separate their most sensitive data from less sensitive, publicly available information on their internal networks. In many environments, internal users have access to all or most of the information on the network. Once attackers have penetrated such a network, they can easily find and exfiltrate important information with little resistance. In several high-profile breaches over the past two years, attackers were able to gain access to sensitive data stored on the same servers with the same level of access as far less important data.

How can this control be implemented, automated, and its effectiveness measured?

1. QW: Organizations should establish a multi-level data identification/separation scheme (e.g., a three- or four-tiered scheme with data separated into categories based on the impact of exposure of the data).
2. QW: Organizations should ensure that files shares have defined controls (such as Windows share access control lists) that specify at least that only "authenticated users" can access the share.
3. Vis/Attrib: Organizations should enforce detailed audit logging for access to non-public data and special authentication for sensitive data.

Associated NIST SP 800-53 Rev 3 Priority 1 Controls:

AC-1, AC-2 (b, c), AC-3 (4), AC-4, AC-6, MP-3, RA-2 (a)

Procedures and tools for implementing this control:

This control is often implemented using the built-in separation of administrator accounts from non-administrator accounts included in most operating systems. While these features are available in most systems, it is important that organizations diligently implement and follow procedures for when administrator-level accounts should be used versus non-administrator accounts.

Control 9 Metric:

The system must be able to detect all attempts by users to access files on local systems or network-accessible file shares without the appropriate privileges and must generate an alert or e-mail for administrative personnel within 24 hours. While the 24 hour timeframe represents the current metric to help organizations improve their current state of security, in the future, organizations should strive for even more rapid alerting, with notification about unauthorized access attempts being sent within two minutes.

Control 9 Test:

To evaluate the implementation of Control 9 on a periodic basis, the evaluation team must create two test accounts each on ten representative systems in the enterprise: five server machines and five client systems. On each system under evaluation, one account must have limited privileges, while the other must have privileges necessary to create files on the systems. The evaluation team must then verify that the non-privileged account is unable to access the files created for the other account on the system. The team must also verify that an alert or e-mail is generated based on the attempted unsuccessful access within 24 hours. At the completion of the test, these accounts must be removed.

Critical Control 10: Continuous Vulnerability Assessment and Remediation

How do attackers exploit the lack of this control?

Soon after new vulnerabilities are discovered and reported by security researchers or vendors, attackers engineer exploit code and then launch that code against targets of interest. Any significant delays in finding or fixing software with critical vulnerabilities provides ample opportunity for persistent attackers to break through, gaining control over the vulnerable machines and getting access to the sensitive data they contain. Organizations that do not scan for vulnerabilities and address discovered flaws proactively face a significant likelihood of having their computer systems compromised.

How can this control be implemented, automated, and its effectiveness measured?

1. QW: Organizations should run automated vulnerability scanning tools against all systems on their networks on a weekly or more frequent basis. Where feasible, vulnerability scanning should occur on a daily basis using an up-to-date vulnerability scanning tool.
2. Config/Hygiene: Organizations should ensure that vulnerability scanning is performed in authenticated mode (i.e., configuring the scanner with administrator credentials) at least quarterly, either with agents running locally on each end system to analyze the security configuration or with remote scanners that are given administrative rights on the system being tested, to overcome limitations of unauthenticated vulnerability scanning.
3. Config/Hygiene: Organizations should compare the results from back-to-back vulnerability scans to verify that vulnerabilities were addressed either by patching, implementing a compensating control, or by documenting and accepting a reasonable business risk. Such acceptance of business risks for existing vulnerabilities should be

periodically reviewed to determine if newer compensating controls or subsequent patches can address vulnerabilities that were previously accepted, or if conditions have changed increasing the risk.

4. Config/Hygiene: Vulnerability scanning tools should be tuned to compare services that are listening on each machine against a list of authorized services. The tools should be further tuned to identify changes over time on systems for both authorized and unauthorized services. Organizations should use government-approved scanning configuration files for their scanning to ensure minimum standards are met.
5. Config/Hygiene: Security personnel should chart the numbers of unmitigated, critical vulnerabilities, for each department/division.
6. Config/Hygiene: Security personnel should share vulnerability reports indicating critical issues with senior management to provide effective incentives for mitigation.
7. Config/Hygiene: Organizations should measure the delay in patching new vulnerabilities and ensure the delay is equal to or less than the benchmarks set forth by the organization, which should be no more than a week for critical patches unless a mitigating control that blocks exploitation is available.
8. Config/Hygiene: Critical patches must be evaluated in a test environment before being pushed into production on enterprise systems. If such patches break critical business applications on test machines, the organization must devise other mitigating controls that block exploitation on systems where the patch cannot be deployed because of its impact on business functionality.
9. Advanced: Organizations should deploy automated patch management tools and software update tools for all systems for which such tools are available and safe.

Associated NIST SP 800-53 Rev 3 Priority 1 Controls:

RA-3 (a, b, c, d), RA-5 (a, b, 1, 2, 5, 6)

Procedures and tools for implementing this control:

A large number of vulnerability scanning tools are available to evaluate the security configuration of systems. Some enterprises have also found commercial services using remotely managed scanning appliances to be effective as well. To help standardize the definitions of discovered vulnerabilities in multiple departments of an agency or even across agencies, it is preferable to use vulnerability scanning tools that measure security flaws and map them to vulnerabilities and issues categorized using one or more of the following industry-recognized vulnerability, configuration, and platform classification schemes and languages: CVE, CCE, OVAL, CPE, CVSS, and/or XCCDF.

Advanced vulnerability scanning tools can be configured with user credentials to login to scanned systems and perform more comprehensive scans than can be achieved without login credentials. For example, organizations can run scanners every week or every month without credentials for an initial inventory of potential vulnerabilities. Then, on a less frequent basis,

such as monthly or quarterly, the organization can run the same scanning tool with user credentials or a different scanning tool that supports scanning with user credentials to find additional vulnerabilities. The frequency of scanning activities, however, should increase as the diversity of an organization's systems increases to account for the varying patch cycles of each vendor.

In addition to the scanning tools that check for vulnerabilities and misconfigurations across the network, various free and commercial tools can evaluate security settings and configurations of local machines on which they are installed. Such tools can provide fine-grained insight into unauthorized changes in configuration or the introduction of security weaknesses inadvertently by administrators.

Effective organizations link their vulnerability scanners with problem-ticketing systems that automatically monitor and report progress on fixing problems and that make visible unmitigated critical vulnerabilities to higher levels of management to ensure the problems are solved.

The most effective vulnerability scanning tools compare the results of the current scan with previous scans to determine how the vulnerabilities in the environment have changed over time. Security personnel use these features to conduct vulnerability trending from month-to-month.

As vulnerabilities related to unpatched systems are discovered by scanning tools, security personnel should determine and document the amount of time that elapsed between the public release of a patch for the system and the occurrence of the vulnerability scan. If this time window exceeds the organization's benchmarks for deployment of the given patch's criticality level, security personnel should note the delay and determine if a deviation was formally documented for the system and its patch. If not, the security team should work with management to improve the patching process.

Additionally, some automated patching tools may not detect or install certain patches, due to error on the vendor's or administrator's part. Because of this, all patch checks should reconcile system patches with a list of patches each vendor has announced on its website.

Control 10 Metric:

All machines identified by the asset inventory system associated with Control 1 must be scanned for vulnerabilities. Additionally, if the vulnerability scanner identifies any devices not included in the asset inventory, it must alert or send email to enterprise administrative personnel within 24 hours. The system must be able to alert or e-mail enterprise administrative personnel within one hour of weekly or daily automated vulnerability scans being completed. If

a scan cannot be completed successfully, the system must alert or send e-mail to administrative personnel within one hour indicating that the scan has not completed successfully. Every 24 hours after that point, the system must alert or send e-mail about the status of uncompleted scans, until normal scanning resumes.

Automated patch management tools must alert or send e-mail to administrative personnel within 24 hours of the successful installation of new patches. While the 24 hour and one hour timeframes represent the current metric to help organizations improve their current state of security, in the future, organizations should strive for even more rapid alerting, with notification about an unauthorized asset connected to the network or an incomplete vulnerability scan being sent within two minutes.

Control 10 Test:

To evaluate the implementation of Control 10 on a periodic basis, the evaluation team must verify that scanning tools have successfully completed their weekly or daily scans for the previous 30 cycles of scanning by reviewing archived alerts and reports to ensure that the scan was completed. If a scan could not be completed in that timeframe, the evaluation team must verify that an alert or e-mail was generated indicating that the scan did not complete.

The evaluation team must verify that application vulnerability scanning tools have successfully completed their regular scans for the previous 30 cycles of scanning by reviewing archived alerts and reports to ensure that the scan was completed. If a scan was not completed successfully, the system must alert or send e-mail to enterprise administrative personnel indicating that the scan did not complete successfully. If a scan could not be completed in that timeframe, the evaluation team must verify that an alert or e-mail was generated indicating that the scan did not complete.

Critical Control 11: Account Monitoring and Control

How do attackers exploit the lack of this control?

Attackers frequently discover and exploit legitimate but inactive user accounts to impersonate legitimate users, thereby making discovery of attacker behavior difficult for network watchers. Accounts of contractors and employees who have been terminated have often been misused in this way. Additionally, some malicious insiders or former employees have accessed accounts left behind in a system long after contract expiration, maintaining their access to an organization's computing system and sensitive data for unauthorized and sometimes malicious purposes.

How can this control be implemented, automated, and its effectiveness measured?

1. QW: Review all system accounts and disable any account that cannot be associated with a business process and business owner.
2. QW: Systems should automatically create a report on a daily basis that includes a list of locked out accounts, disabled accounts, accounts with passwords that exceed the maximum password age, and accounts with passwords that never expire. This list should be sent to the associated system administrator in a secure fashion.
3. QW: Organizations should establish and follow a process for revoking system access by disabling accounts immediately upon termination of an employee or contractor.
4. QW: Organizations should regularly monitor the use of all accounts, automatically logging off users after a standard period of inactivity.
5. QW: Organizations should monitor account usage to determine dormant accounts that have not been used for a given period, such as 30 days, notifying the user or user's manager of the dormancy. After a longer period, such as 60 days, the account should be disabled.
6. QW: On a periodic basis, such as quarterly or at least annually, organizations should require that managers match active employees and contractors with each account belonging to their managed staff. Security or system administrators should then disable accounts that are not assigned to active employees or contractors.
7. QW: When a dormant account is disabled, any files associated with that account should be encrypted and moved to a secure file server for analysis by security or management personnel.
8. Vis/Attrib: Organizations should monitor attempts to access deactivated accounts through audit logging.
9. Config/Hygiene: Organizations should profile each user's typical account usage by determining normal time-of-day access and access duration for each user. Daily reports should be generated that indicate users who have logged in during unusual hours or have exceeded their normal login duration by 150%.

Associated NIST SP 800-53 Rev 3 Priority 1 Controls:

AC-2 (e, f, g, h, j, 2, 3, 4, 5), AC-3

Procedures and tools for implementing this control:

Although most operating systems include capabilities for logging information about account usage, these features are sometimes disabled by default. Even when such features are present and active, they often do not provide fine-grained detail about access to the system by default. Security personnel can configure systems to record more detailed information about account access, and utilize home-grown scripts or third-party log analysis tools to analyze this information and profile user access of various systems.

Control 11 Metric:

The system must be capable of identifying unauthorized user accounts when they exist on the system. An automated list of user accounts on the system must be created every 24 hours and an alert or e-mail must be sent to administrative personnel within one hour of completion of a list being created. While the one-hour timeframe represents the current metric to help organizations improve their current state of security, in the future, organizations should strive for even more rapid alerting, with notification regarding the creation of the list of user accounts being sent within two minutes.

Control 11 Test:

To evaluate the implementation of Control 11 on a periodic basis, the evaluation team must verify that the list of locked out accounts, disabled accounts, accounts with passwords that exceed the maximum password age, and accounts with passwords that never expire has successfully been completed on a daily basis for the previous thirty days by reviewing archived alerts and reports to ensure that the lists were completed. In addition, a comparison of a baseline of allowed accounts must be compared to the accounts that are active in all systems. The report of all differences must be created based on this comparison.

Critical Control 12: Malware Defenses

How do attackers exploit the lack of this control?

Malicious software is an integral and dangerous aspect of Internet threats, targeting end-users and organizations via web browsing, email attachments, mobile devices, and other vectors. Malicious code may tamper with the system's contents, capture sensitive data, and spread to other systems. Modern malware aims to avoid signature-based and behavioral detection, and may disable anti-virus tools running on the targeted system. Anti-virus and anti-spyware software, collectively referred to as anti-malware tools, help defend against these threats by attempting to detect malware and block its execution.

How can this control be implemented, automated, and its effectiveness measured?

1. QW: Organizations should monitor workstations, servers, and mobile devices for active, up-to-date anti-malware protection with anti-virus, anti-spyware, and host-based Intrusion Prevention System functionality. Enterprise administrative features should be used to check daily the number of systems that do not have the latest anti-malware

signatures, keeping the number of such systems small or eliminating them entirely through rapid and continuous updates. All malware detection events should be sent to enterprise anti-malware administration tools and event log servers.

2. QW: Organizations should employ anti-malware software and signature auto update features or have administrators manually push updates to all machines on a daily basis. After applying an update, automated systems should verify that each system has received its signature update.
3. QW: Organizations should configure laptops, workstations, and servers so that they will not auto-run content from USB tokens (i.e., "thumb drives"), USB hard drives, CDs/DVDs, Firewire devices, external SATA devices, mounted network shares, or other removable media.
4. QW: Organizations should configure systems so that they conduct an automated anti-malware scan of removable media when it is inserted.
5. Advanced: Organizations should deploy honeypots or tarpits as detection mechanisms that can also slow down an attacker's progress inside a network.
6. Advanced: Organizations should deploy Network Access Control (NAC) tools to verify security configuration and patch level compliance before granting access to a network.

Associated NIST SP 800-53 Rev 3 Priority 1 Controls:

SC-18, SC-26, SI-3 (a, b, 1, 2, 5, 6)

Procedures and tools for implementing this control:

Relying on policy and user action to keep anti-malware tools up to date has been widely discredited, as many users have not proven able to keep such tools up to date consistently. To ensure anti-virus signatures are up to date, effective organizations use automation. They use the built-in administrative features of enterprise end-point security suites to verify that anti-virus, anti-spyware, and host-based IDS features are active on every managed system. They run automated assessments daily and review the results to find and mitigate systems that have deactivated such protections, as well as systems that do not have the latest malware definitions. For added security in depth, and for those systems that may fall outside the enterprise anti-malware coverage, some organizations use network access control technology that tests machines for compliance with security policy before allowing them to connect to the network.

Some enterprises deploy free or commercial honeypot and tarpit tools to identify attackers in their environment. Security personnel should continuously monitor honeypots and tarpits to determine whether traffic is directed to them and account logins are attempted. When they identify such events, these personnel should gather the source address from which this traffic originates and other details associated with the attack for a follow-on investigation.

Control 12 Metric:

The system must identify any malicious software that is installed, attempted to be installed, executed, or attempted to be executed on a computer system within one hour, alerting or sending email notification to a list of enterprise personnel via their centralized anti-malware console or event log system. Systems must block installation, prevent execution, or quarantine malicious software within one 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 malicious code until such time as the threat has been completely mitigated on that system. While the one hour timeframe represents the current metric to help organizations improve their current state of security, in the future, organizations should strive for even more rapid detection and malware isolation, with notification about malware in the enterprise being sent within two minutes and blocking, execution prevention, or quarantine actions occurring within five minutes.

Control 12 Test:

To evaluate the implementation of Control 12 on a periodic basis, the evaluation team must move a benign software test program which appears to be malware (such as an EICAR file or benign hacker tools) that is not included in the official authorized software list to ten systems on the network via a network share. The selection of these systems must be as random as possible and include a cross-section of the organization's systems and locations. The evaluation team must then verify that the systems generate an alert or e-mail notice regarding the benign malware within one hour. The team must also verify that the alert or e-mail is received within one hour indicating that the software has been blocked or quarantined. The evaluation team must verify that the system provides details of the location of each machine with this new test file, including information about the asset owner. The evaluation team must then verify that the file is blocked by attempting to execute or open it and verifying that it is not allowed to be accessed.

Once this test has been performed transferring the files to organization systems via removable media, the same test must be repeated, but transferring the benign malware to ten systems via e-mail instead. The organization must expect the same notification results as noted with the removable media test.

Critical Control 13: Limitation and Control of Network Ports, Protocols, and Services

How do attackers exploit the lack of this control?

Attackers search for remotely accessible network services that are vulnerable to exploitation. Common examples include poorly configured web servers, mail servers, file and print services, and DNS servers installed by default on a variety of different device types, often without a business need for the given service. Many software packages automatically install services and turn them on as part of the installation of the main software package without informing a user or administrator that the services have been enabled. Attackers scan for such issues and attempt to exploit these services, often attempting default user IDs and passwords or widely available exploitation code.

How can this control be implemented, automated, and its effectiveness measured?

1. QW: Host-based firewalls or port filtering tools should be applied on end systems, with a default-deny rule that drops all traffic except those services and ports that are explicitly allowed.
2. Config/Hygiene: Services needed for business use across the internal network should be reviewed quarterly via a change control group, and business units should re-justify the business use. Sometime services are turned on for projects or limited engagements, and should be turned off when they are no longer needed.
3. Config/Hygiene: Operate critical services on separate physical host machines, such as DNS, file, mail, web, and database servers.

Associated NIST SP 800-53 Rev 3 Priority 1 Controls:

CM-6 (a, b, d, 2, 3), CM-7 (1), SC-7 (4, 5, 11, 12)

Procedures and tools for implementing this control:

Port scanning tools are used to determine which services are listening on the network for a range of target systems. In addition to determining which ports are open, effective port scanners can be configured to identify the version of the protocol and service listening on each discovered open port. This list of services and their versions are compared against an inventory of services required by the organization for each server and workstation, in an asset management system, such as those described in Critical Control 1. Recently added features in these port scanners are being used to determine the changes in services offered by scanned machines on the network since the previous scan, helping security personnel identify differences over time.

Control 13 Metric:

The system must be capable of identifying any new unauthorized listening network ports that are connected to the network within 24 hours, alerting or sending email notification to a list of enterprise personnel. Every 24 hours after that point, the system must alert or send e-mail

about the status of the system until the listening network port has been disabled or it has been authorized by change management. The system service baseline database and alerting system must be able to identify the location, department, and other details about the system where authorized and unauthorized network ports are running. While the 24 hour timeframe represents the current metric to help organizations improve their current state of security, in the future, organizations should strive for even more rapid alerting, with notification about an unauthorized open port on the network being sent within two minutes.

Control 13 Test:

To evaluate the implementation of Control 13 on a periodic basis, the evaluation team must install hardened test services with network listeners on ten locations on the network, including a selection of subnets associated with DMZs, workstations, and servers. The selection of these systems must be as random as possible and include a cross-section of the organization's systems and locations. The evaluation team must then verify that the systems generate an alert or e-mail notice regarding the newly installed services within 24 hours of the services being installed on the network. The test team must verify that the system provides details of the location of all of the systems where test services have been installed.

Critical Control 14: Wireless Device Control

How do attackers exploit the lack of this control?

Major data thefts have been initiated by attackers who have gained wireless access to organizations from nearby parking lots, bypassing organizations' security perimeters by connecting wirelessly to access points inside the organization. Wireless clients accompanying travelling officials are infected on a regular basis through remote exploitation during air travel or in cyber cafés. Such exploited systems are then used as back doors when they are reconnected to the network of a target organization. Still other organizations have reported the discovery of unauthorized wireless access points on their networks, planted and sometimes hidden for unrestricted access to an internal network. Because they do not require direct physical connections, wireless devices are a convenient vector for attackers to maintain long-term access into a target environment.

How can this control be implemented, automated, and its effectiveness measured?

1. QW: Organizations should ensure that each wireless device connected to the network matches an authorized configuration and security profile, with a documented owner of

the connection and a defined business need. Organizations should deny access to those wireless devices that do not have such a configuration and profile.

2. QW: Organizations should ensure that all wireless access points are manageable using enterprise management tools. Access points designed for home use often lack such enterprise management capabilities, and should therefore be avoided in enterprise environments.
3. QW: Network vulnerability scanning tools should be configured to detect wireless access points connected to the wired network. Identified devices should be reconciled against a list of authorized wireless access points. Unauthorized (i.e., rogue) access points should be deactivated.
4. Vis/Attrib: Organizations should use wireless intrusion detection systems (WIDS) to identify rogue wireless devices and detect attack attempts and successful compromise. In addition to WIDS, all wireless traffic should be monitored by a wired IDS as traffic passes into the wired network.
5. Config/Hygiene: Where a specific business need for wireless access has been identified, organizations should configure wireless access on client machines to allow access only to authorized wireless networks.
6. Config/Hygiene: For devices that do not have an essential wireless business purpose, organizations should disable wireless access in the hardware configuration (BIOS or EFI), with password protections to lower the possibility that the user will override such configurations.
7. Config/Hygiene: Organizations should regularly scan for unauthorized or misconfigured wireless infrastructure devices, using techniques such as “war driving” to identify access points and clients accepting peer-to-peer connections. Such unauthorized or misconfigured devices should be removed from the network, or have their configurations altered so that they comply with the security requirements of the organization.
8. Config/Hygiene: Organizations should ensure all wireless traffic leverages at least AES encryption used with at least WPA2 protection.
9. Config/Hygiene: Organizations should ensure wireless networks use authentication protocols such as EAP/TLS or PEAP, which provide credential protection and mutual authentication.
10. Config/Hygiene: Organizations should ensure wireless clients use strong, multi-factor authentication credentials to mitigate the risk of unauthorized access from compromised credentials.
11. Config/Hygiene: Organizations should disable peer-to-peer wireless network capabilities on wireless clients, unless such functionality meets a documented business need.
12. Config/Hygiene: Organizations should disable wireless peripheral access of devices (such as Bluetooth), unless such access is required for a documented business need.
13. Advanced: Organizations should configure all wireless clients used to access agency networks or handle organization data in a manner so that they cannot be used to connect to public wireless networks or any other networks beyond those specifically allowed by the agency.

Associated NIST SP 800-53 Rev 3 Priority 1 Controls:

AC-17, AC-18 (1, 2, 3, 4), SC-9 (1), SC-24, SI-4 (14, 15)

Procedures and tools for implementing this control:

Effective organizations run commercial wireless scanning, detection, and discovery tools as well as commercial wireless intrusion detection systems. Additionally, the security team could periodically capture wireless traffic from within the borders of a facility and use free and commercial analysis tools to determine whether the wireless traffic was transmitted using weaker protocols or encryption than the organization mandates. When devices relying on weak wireless security settings are identified, they should be found within the organization's asset inventory and either reconfigured more securely or denied access to the agency network.

Additionally, the security testing team could employ remote management tools on the wired network to pull information about the wireless capabilities and devices connected to managed systems.

Control 14 Metric:

The system must be capable of identifying unauthorized wireless devices or configurations when they are within range of the organization's systems or connected to their networks. The system must be capable of identifying any new unauthorized wireless devices that associate or join the network within one hour, alerting or sending email notification to a list of enterprise personnel. The system must automatically isolate an attached wireless access point from the network within one hour and alert or send e-mail notification when isolation is achieved. 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. The asset inventory database and alerting system must be able to identify the location, department, and other details of where authorized and unauthorized wireless devices are plugged into the network. While the 24 hour and one hour timeframes represent the current metric to help organizations improve their current state of security, in the future, organizations should strive for even more rapid alerting and isolation, with notification about an unauthorized wireless devices being sent within two minutes and isolation within five minutes.

Control 14 Test:

To evaluate the implementation of Control 14 on a periodic basis, the evaluation team staff must configure ten unauthorized but hardened wireless clients and wireless access points to the organization's network and attempt to connect them to the organization's wireless

networks. In the case of wireless access points, these access points must not be directly connected to the organization's trusted network. Instead they must simply be configured to act as a wireless gateway without physically connecting to a wired network interface. In the case of scanning for wireless access points from a wired interface, the connected access point must have the wireless radio disabled for the duration of the test. These systems must be configured to test each of the following scenarios:

- A wireless client with an unauthorized SSID configured on it.
- A wireless client with improper encryption configured.
- A wireless client with improper authentication configured.
- A wireless access point with improper encryption configured.
- A wireless access point with improper authentication configured.
- A completely rogue wireless access point using an unauthorized configuration.

When each of the above noted systems are attempting to connect to the wireless network, an alert must be generated and enterprise staff must respond to the alerts to isolate the detected device or remove the device from the network.

Critical Control 15: Data Loss Prevention

How do attackers exploit the lack of this control?

In recent years, attackers have exfiltrated more than 20 terabytes of often sensitive data from Department of Defense and Defense Industrial Base organizations (e.g., contractors doing business with the DoD), as well as civilian government organizations. Many attacks occurred across the network, while others involved physical theft of laptops and other equipment holding sensitive information. Yet, in most cases, the victims were not aware that significant amounts of sensitive data were leaving their systems because they were not monitoring data outflows. The movement of data across network boundaries both electronically and physically must be carefully scrutinized to minimize its exposure to attackers.

The loss of control over protected or sensitive data by organizations is a serious threat to business operations as well as, potentially, national security. While some data is leaked or lost as a result of theft or espionage, the vast majority of these problems result from poorly understood data practices, a lack of effective policy architectures, and user error. Data loss can even occur as a result of legitimate activities such as e-Discovery during litigation, particularly when records retention practices are ineffective or non-existent.

The phrase “Data Loss Prevention” (DLP) refers to a comprehensive approach covering people, processes, and systems that identify, monitor, and protect data in use (e.g., endpoint actions), data in motion (e.g., network actions), and data at rest (e.g., data storage) through deep content inspection and with a centralized management framework. Over the last several years, there has been a noticeable shift in attention and investment from securing the network, to securing systems within the network, to securing the data itself. DLP controls are based on policy, and include classifying sensitive data, discovering that data across an enterprise, enforcing controls, and reporting and auditing to ensure policy compliance.

How can this control be implemented, automated, and its effectiveness measured?

1. QW: Organizations should deploy approved hard drive encryption software to laptop machines that hold sensitive data.
2. Vis/Attrib: Network monitoring tools should analyze outbound traffic looking for a variety of anomalies, including large file transfers, long-time persistent connections, connections at regular repeated intervals, unusual protocols and ports in use, and possibly the presence of certain keywords in the data traversing the network perimeter.
3. Vis/Attrib: Deploy an automated tool on network perimeters that monitors for certain Personally Identifiable Information (PII), keywords, and other document characteristics in an automated fashion to determine attempts to exfiltrate data in an unauthorized fashion across network boundaries and block such transfers while alerting information security personnel.
4. Vis/Attrib: Conduct periodic scans of server machines using automated tools to determine whether PII data is present on the system in clear text. These tools, which search for patterns that indicate the presence of PII, can help identify if a business or technical process is leaving behind or otherwise leaking sensitive information in data at rest.
5. Config/Hygiene: Data should be moved between networks using secure, authenticated, encrypted mechanisms.
6. Config/Hygiene: Data stored on removable, easily transported storage media, such as USB tokens (i.e., “thumb drives”), USB portable hard drives, and CDs/DVDs, should be encrypted. Systems should be configured so that all data written to such media is automatically encrypted without user intervention.
7. Advanced: If there is no business need for supporting such devices, organizations should configure systems so that they will not write data to USB tokens or USB hard drives. If such devices are required, utilize enterprise software that can configure systems to allow only specific USB devices (based on serial number or other unique property) to be accessed, and that all data placed on such devices be automatically encrypted.

Associated NIST SP 800-53 Rev 3 Priority 1 Controls:

AC-4, MP-2 (2), MP-4 (1), SC-7 (6, 10), SC-9, SC-13, SC-28 (1), SI-4 (4, 11), PM-7

Procedures and tools for implementing this control:

Commercial DLP solutions are available to look for exfiltration attempts and detect other suspicious activities associated with a protected network holding sensitive information. Organizations deploying such tools should carefully inspect their logs and follow-up on any discovered attempts, even those that are successfully blocked, to transmit sensitive information out of the organization without authorization.

Control 15 Metric:

The system must be capable of identifying unauthorized data leaving the organization's systems whether via network file transfers or removable media. Within one hour of a data exflitration event or attempt taking place, enterprise administrative personnel must be alerted by the appropriate monitoring system. Once the alert has been generated it must also note the system and location where the event or attempt occurred. If the system is in the organization's asset management database, then the system owner must also be indicated in the generated alerts. Every 24 hours after that point, the system must alert or send e-mail about the status of the systems until the source of the event has been identified and the risk mitigated. While the one hour timeframe represents the current metric to help organizations improve their current state of security, in the future, organizations should strive for even more rapid alerting, with notification about data exfiltration events or attempts being sent within two minutes.

Control 15 Test:

To evaluate the implementation of Control 15 on a periodic basis, the evaluation team must attempt to move test data sets (that trigger DLP systems but do not contain sensitive data) outside of the trusted computing environment via both network file transfers and via removable media. Each of the following tests must be performed at least three times:

- Attempting to transfer large data sets across network boundaries from an internal system.
- Attempting to transfer test data sets of PII (that trigger DLP systems but do not contain sensitive data) across network boundaries from an internal system (using multiple keywords specific to the business).
- Attempting to maintain a persistent network connection for at least ten hours across network boundaries between an internal and external system, although little data may be exchanged.
- Attempting to maintain a network connection across network boundaries using an anomalous service port number between an internal and external system.
- Inserting a USB token into an organization system and attempting to transfer example test data to the USB device.

Each of these tests must be performed from multiple, widely distributed systems on the organization's network in order to test the effectiveness of the monitoring systems. Once each of these events has occurred, the time it takes for enterprise staff to respond to the event must be recorded.

Additional Controls

The following sections identify additional controls that are important but cannot be fully automatically or continuously monitored to the same degree as the controls covered earlier in this document.

Critical Control 16: Secure Network Engineering

How do attackers exploit the lack of this control?

Many controls in this document are effective but can be circumvented in networks that are poorly designed. Without a carefully planned and properly implemented network architecture, attackers can bypass security controls on certain systems, pivoting through the network to gain access to target machines. Attackers frequently map networks looking for unneeded connections between systems, weak filtering, and a lack of network separation. Therefore a robust, secure network engineering process must be employed to complement the detailed controls being measured in other sections of this document.

How can this control be implemented and its effectiveness measured?

Among the engineering/architectural standards to be used are:

1. QW: Each organization should standardize the DHCP lease information and time assigned to systems, and verbosely log all information about DHCP leases distributed in the organization.

2. Config/Hygiene: To support rapid response and shunning of detected attacks, the network architecture and the systems that make it up should be engineered for rapid deployment of new access control lists, rules, signatures, blocks, blackholes, and other defensive measures.

3. Vis/Attrib: DNS should be deployed in a hierarchical, structured fashion, with all internal network client machines configured to send requests to intranet DNS servers, not to DNS servers located on the Internet. These internal DNS servers should be configured to forward requests they cannot resolve to DNS servers located on a protected DMZ. These DMZ servers, in turn, should be the only DNS servers allowed to send requests to the Internet.

4. Advanced: Organizations should segment the enterprise network into multiple, separate trust zones to provide more granular control of system access and additional intranet boundary defenses.

Associated NIST SP 800-53 Rev 3 Priority 1 Controls:

IR-4 (2), SA-8, SC-7 (1, 13), SC-20, SC-21, SC-22, PM-7

Procedures and tools for implementing this control:

To help ensure a consistent, defensible network, the architecture of each network should be based on a template that describes the overall layout of the network and the services it provides. Organizations should prepare network diagrams for each of their networks that show network components such as routers, firewalls, and switches, along with significant servers and groups of client machines.

Critical Control 17: Penetration Tests and Red Team Exercises

How do attackers exploit the lack of this control?

Attackers penetrate networks and systems through social engineering and by exploiting vulnerable software and hardware. Once they get access, they often burrow deep into target systems and broadly expand the number of machines over which they have control. Most organizations do not exercise their defenses so they are uncertain about their capabilities and unprepared for identifying and responding to attack.

Penetration testing involves mimicking the actions of computer attackers to identify vulnerabilities in a target organization, and exploiting them to determine what kind of access an

attacker can gain. Penetration tests typically provide a deeper analysis of security flaws than the vulnerability assessments described in Control 10. Vulnerability assessments focus on identifying potential vulnerabilities, while penetration testing goes deeper with controlled attempts at exploiting vulnerabilities, approaching target systems as an attacker would. The result provides deeper insight into the business risks of various vulnerabilities, by showing whether and how an attacker can compromise machines, pivot to other systems inside a target organization, and gain access to sensitive information assets.

Red team exercises go further than penetration testing. Red team exercises have the goals of improved readiness of the organization, better training for defensive practitioners, and inspection of current performance levels. Independent red teams can provide valuable objectivity regarding both the existence of vulnerabilities and the efficacy of defenses and mitigating controls already in place and even those planned for future implementation.

How can this control be implemented and its effectiveness measured?

1. QW: Organizations should conduct regular penetration tests to identify vulnerabilities and attack vectors that can be used to exploit enterprise systems successfully. Penetration testing should occur from outside the network perimeter (i.e., the Internet or wireless frequencies around an organization) as well from within its boundaries (i.e., on the internal network) to simulate both outsider and insider attacks.
2. Vis/Attrib: Organizations should perform periodic red team exercises to test the readiness of organizations to identify and stop attacks or to respond quickly and effectively.
3. Vis/Attrib: Organizations should ensure systemic problems discovered in penetration tests and red team exercises are fully mitigated.
4. Vis/Attrib: Organizations should measure how well the organization has reduced the significant enablers for attackers by setting up automated processes to find:
   - Cleartext emails and documents with “password” in the filename or body.
   - Critical network diagrams stored online and in cleartext
   - Critical configuration files stored online and in cleartext.
   - Vulnerability assessment, penetration test reports, and red team findings documents stored online and in cleartext.
   - Other sensitive information identified by management personnel as critical to the operation of the enterprise during the scoping of a penetration test or red team exercise.
5. Advanced: Organizations should devise a scoring method for determining the results of red team exercises so that results can be compared over time.
6. Advanced: Organizations should create a test bed that mimics a production environment for specific penetration tests and red team attacks against elements that are not typically tested in production, such as attacks against SCADA and other control systems.

Associated NIST SP 800-53 Rev 3 Priority 1 Controls:

CA-2 (1, 2), CA-7 (1, 2), RA-3, RA-5 (4, 9), SA-12 (7)

Procedures and tools for implementing this control:

Each organization should define a clear scope and rules of engagement for penetration testing and red team analyses. The scope of such projects should include, at least, systems with the highest value information and production processing functionality of the organization. Other, lowered value systems may also be tested to see if they can be used as pivot points to compromise higher-valued targets. The rules of engagement for penetration tests and red team analyses should describe, at a minimum, times of day for testing, duration of tests, and overall test approach.

Critical Control 18: Incident Response Capability

How do attackers exploit the lack of this control?

A great deal of damage has been done to organizational reputations and a great deal of information has been lost in organizations that do not have fully effective incident response programs in place. Without an incident response plan, an organization may not discover an attack in the first place, or, if the attack is detected, the organization may not follow proper procedures to contain damage, eradicate the attacker's presence, and recover in a secure fashion. Thus, the attacker may have far higher impact on the target organization, causing more damage, infecting more systems, and possibly exfiltrating more sensitive data than would otherwise be possible with an effective incident response plan.

The National Institute of Standards and Technology (NIST) has released detailed guidelines for creating and running an incident response team in Special Publication 800-61, available at http://csrc.nist.gov/publications/nistpubs/800-61/sp800-61.pdf.

How can this control be implemented and its effectiveness measured?

Among the most important elements included in these guidelines are:

1. QW: Organizations should ensure that they have written incident response procedures, which include a definition of personnel roles for handling incidents. The procedures should define the phases of incident handling consistent with the NIST guidelines cited above.
2. QW: Organizations should assign job titles and duties for handling computer and network incidents to specific individuals.

3. QW: Organizations should define management personnel that will support the incident handling process within each organization, acting in key decision-making roles.
4. QW: Organizations should devise organization-wide standards for the time required for system administrators and other personnel to report anomalous events to the agency incident handling team, the mechanisms for such reporting, and the kind of information that should be passed in the incident notification. This reporting should also include notifying US-CERT in accordance with federal requirements for involving that organization in computer incidents.
5. QW: Organizations should publish information to all personnel, including employees and contractors, regarding reporting computer anomalies and incidents to the incident handling team. Include such information in routine employee awareness activities.
6. Config/Hygiene: Organizations should conduct periodic incident scenario sessions for personnel associated with the incident handling team to ensure that personnel understand current threats and risks, as well as their responsibilities in supporting the incident handling team.

Associated NIST SP 800-53 Rev 3 Priority 1 Controls:

IR-1, IR-2 (1), IR-4, IR-5, IR-6 (a), IR-8

Procedures and tools for implementing this control:

After defining detailed incident response procedures, the incident response team should engage in periodic scenario-based training, working through a series of attack scenarios fine-tuned to the threats and vulnerabilities the organization faces. These scenarios help ensure that team members understand their role on the incident response team and also help prepare them to handle incidents.

Critical Control 19: Data Recovery Capability

How do attackers exploit the lack of this control?

When attackers compromise machines, they often make significant changes to configurations and software. Sometimes attackers also make subtle alterations of data stored on compromised machines, potentially jeopardizing organizational effectiveness with polluted information. When the attackers' presence is discovered, organizations without a trustworthy data recovery capability can have extreme difficulty removing all aspects of the attacker's presence on the machine.

How can this control be implemented and its effectiveness measured?

1. QW: Organizations should ensure that each system is automatically backed up on at least a weekly basis, and more often for systems storing sensitive information. To help ensure the ability to rapidly restore a system from backup, make sure that the operating system, application software, and data on a machine are each included in the overall backup procedure. These three components of a system do not have to be included in the same backup file or using the same backup software. However, each must be backed up at least weekly.
2. Config/Hygiene: Organizations should ensure that backups are encrypted when they are stored locally, as well as when they are moved across the network.
3. Config/Hygiene: Backup media, such as hard drives and tapes, should be stored in physically secure, locked facilities.

Associated NIST SP 800-53 Rev 3 Priority 1 Controls:

CP-9 (a, b, d, 1, 3), CP-10 (6)

Procedures and tools for implementing this control:

Once per quarter, a testing team should evaluate a random sample of system backups by attempting to restore them on a test bed environment. The restored systems should be verified to ensure that the operating system, application, and data from the backup are all intact and functional.

Critical Control 20: Security Skills Assessment and Appropriate Training to Fill Gaps

How do attackers exploit the lack of this control?

Five groups of people are constantly being tested by exploitation attempts by attackers:

1. End users are fooled via social engineering scams, in which they are tricked into providing passwords, opening attachments, loading software from untrusted sites, or visiting malicious web sites.
2. System administrators are also fooled in the same manner as normal users but are also tested when attackers attempt to trick the administrator into setting up unauthorized accounts.
3. Security operators and analysts are tested with new and innovative attacks introduced on a continual basis.
4. Application programmers are tested by criminals who find and exploit the vulnerabilities in the code that they write.

5. To a lesser degree, system owners are tested when they are asked to invest in cyber security but are unaware of the devastating impact a compromise and data exfiltration or data alteration would have on their mission.

Any organization that hopes to be ready to find and respond to attacks effectively owes it to their employees and contractors to find the gaps in their knowledge and to provide exercises and training to fill those gaps. A solid security skills assessment program can provide actionable information to decision makers about where security awareness needs to be improved, and can also help determine proper allocation of limited resources to improve security practices.

How can this control be implemented and its effectiveness measured?

1. QW: Organizations should develop security awareness training for various personnel job descriptions. The training should include specific, incident-based scenarios showing the threats an organization faces. The training should reflect proven defenses for the latest attack techniques.
2. Config/Hygiene: Organizations should devise periodic security awareness assessment quizzes, to be given to employees and contractors on at least an annual basis, determining whether they understand the information security policies and procedures for the organization, as well as their role in those procedures.
3. Config/Hygiene: Organizations should conduct periodic exercises to verify that employees and contractors are fulfilling their information security duties, by conducting tests to see whether employees will click on a link from suspicious e-mail or provide sensitive information on the telephone without following appropriate procedures for authenticating a caller.

Associated NIST SP 800-53 Rev 3 Priority 1 Controls:

AT-1, AT-2 (1), AT-3 (1)

Procedures and tools for implementing this control:

The key to upgrading skills is measurement – not with certification examinations, but with assessments that show both the employee and the employer where knowledge is sufficient and where the gaps are. Once the gaps have been identified, those employees who have the requisite skills and knowledge can be called upon to mentor the employees who need skills improvement or the organization can develop training programs that directly fill the gaps and maintain employee readiness.

SUMMARY

This document has been developed through the collaboration of a diverse set of security experts. While there is no such thing as absolute protection, proper implementation of the security controls identified in this document will ensure that an organization is protecting against the most significant attacks. As attacks change, additional controls or tools become available, or the state of common security practice advances, this document will be updated to reflect what is viewed by the collaborating authors as the most important security controls to defend against cyber attacks.

Appendix A: Mapping between Top 20 Critical Controls and NIST SP 800-53 Rev 3 Priority 1 Items

This mapping relates the controls set forth in this document to NIST SP 800-53 Rev 3. Please note that the NIST controls may impose additional requirements beyond those explicitly stated in this document.

Control	References
Critical Control 1: Inventory of Authorized and Unauthorized Devices	CM-8 (a, c, d, 2, 3, 4), PM-5, PM-6
Critical Control 2: Inventory of Authorized and Unauthorized Software	CM-1, CM-2 (2, 4, 5), CM-3, CM-5 (2, 7), CM-7 (1, 2), CM-8 (1, 2, 3, 4, 6), CM-9, PM-6, SA-6, SA-7
Critical Control 3: Secure Configurations for Hardware and Software	CM-1, CM-2 (1, 2), CM-3 (b, c, d, e, 2, 3), CM-5 (2), CM-6 (1, 2, 4), CM-7 (1), SA-1 (a), SA-4 (5), SI-7 (3), PM-6
Critical Control 4: Secure Configurations for Network Devices such as Firewalls, Routers, and Switches	AC-4 (7, 10, 11, 16), CM-1, CM-2 (1), CM-3 (2), CM-5 (1, 2, 5), CM-6 (4), CM-7 (1, 3), IA-2 (1, 6), IA-5, IA-8, RA-5, SC-7 (2, 4, 5, 6, 8, 11, 13, 14, 18), SC-9
Critical Control 5: Boundary Defense	AC-17 (1), AC-20, CA-3, IA-2 (1, 2), IA-8, RA-5, SC-7 (1, 2, 3, 8, 10, 11, 14), SC-18, SI-4 (c, 1, 4, 5, 11), PM-7
Critical Control 6: Maintenance, Monitoring and Analysis of Security Audit Logs	AC-17 (1), AC-19, AU-2 (4), AU-3 (1,2), AU-4, AU-5, AU-6 (a, 1, 5), AU-8, AU-9 (1, 2), AU-12 (2), SI-4 (8)
Critical Control 7: Application Software Security	CM-7, RA-5 (a, 1), SA-3, SA-4 (3), SA-8, SI-3, SI-10
Critical Control 8: Controlled Use of Administrative Privileges	AC-6 (2, 5), AC-17 (3), AC-19, AU-2 (4)
Critical Control 9: Controlled Access Based	AC-1, AC-2 (b, c), AC-3 (4), AC-4, AC-6, MP-3, RA-2 (a)

on Need to Know	
Critical Control 10: Continuous Vulnerability Assessment and Remediation	RA-3 (a, b, c, d), RA-5 (a, b, 1, 2, 5, 6)
Critical Control 11: Account Monitoring and Control	AC-2 (e, f, g, h, j, 2, 3, 4, 5), AC-3
Critical Control 12: Malware Defenses	SC-18, SC-26, SI-3 (a, b, 1, 2, 5, 6)
Critical Control 13: Limitation and Control of Network Ports, Protocols, and Services	CM-6 (a, b, d, 2, 3), CM-7 (1), SC-7 (4, 5, 11, 12)
Critical Control 14: Wireless Device Control	AC-17, AC-18 (1, 2, 3, 4), SC-9 (1), SC-24, SI-4 (14, 15)
Critical Control 15: Data Loss Prevention	AC-4, MP-2 (2), MP-4 (1), SC-7 (6, 10), SC-9, SC-13, SC-28 (1), SI-4 (4, 11), PM-7
Critical Control 16: Secure Network Engineering	IR-4 (2), SA-8, SC-7 (1, 13), SC-20, SC-21, SC-22, PM-7,
Critical Control 17: Penetration Tests and Red Team Exercises	CA-2 (1, 2), CA-7 (1, 2), RA-3, RA-5 (4, 9), SA-12 (7)
Critical Control 18: Incident Response Capability	IR-1, IR-2 (1), IR-4, IR-5, IR-6 (a), IR-8
Critical Control 19: Data Recovery Capability	CP-9 (a, b, d, 1, 3), CP-10 (6)
Critical Control 20: Security Skills Assessment and Appropriate Training To Fill Gaps	AT-1, AT-2 (1), AT-3 (1)

Appendix B: Attack Types

As described in the introduction to the Twenty Critical Controls, numerous contributors who are responsible for responding to actual attacks or conducting Red Team exercises were involved in the creation of this document. The resulting controls are therefore based on first-hand knowledge or real-world attacks and the associated defenses.

Attack Summary	Most Directly Related Control
Attackers continually scan for new, unprotected systems, including test or experimental systems, and exploit such systems to gain control of them.	1
Attackers continually scan for vulnerable software and exploit it to gain control of target machines.	2
Attackers distribute hostile content on Internet-accessible (and sometimes internal) websites that exploits unpatched and improperly secured client software running on victim machines.	2
Attackers use currently infected or compromised machines to identify and exploit other vulnerable machines across an internal network.	2
Attackers exploit weak default configurations of systems that are more geared to ease of use than security.	3
Attackers exploit and infiltrate through network devices whose security configuration has been weakened over time by granting, for specific short-term business needs, supposedly temporary exceptions that are never removed.	4
Attackers exploit boundary systems on Internet-accessible DMZ networks, and then pivot to gain deeper access on internal networks.	5
Attackers operate undetected for extended periods of time on compromised systems because of a lack of logging and log review.	6
Attackers exploit weak application software, particularly web applications, through attack vectors such as SQL injection, Cross-Site-Scripting, and related issues.	7
Attackers trick a user with an administrator level account into opening a phishing-style e-mail with attachment or surfing to the attacker's content	8

on an Internet website, allowing the attacker's malicious code or exploit to run on the victim machine with full administrator privileges.	
Attackers escalate their privileges on victim machines by launching password guessing, password cracking, or privilege escalation exploits to gain administrator control of systems, which is then used to propagate to other victim machines across an enterprise.	8
Attackers gain access to sensitive documents in an organization that does not properly identify and protect sensitive information, separating it from non-sensitive information.	9
Attackers exploit new vulnerabilities on systems that lack critical patches in organizations that do not know that they are vulnerable because they lack continuous vulnerability assessments and effective remediation.	10
Attackers compromised inactive user accounts left behind by temporary workers, contractors, and former employees, including accounts left behind by the attackers themselves who are former employees.	11
Attackers use malicious code to gain and maintain control of target machines, capture sensitive data, and spread to other systems, sometimes wielding code that disables or dodges signature-based anti-virus tools.	12
Attackers scan for remotely accessible services on target systems that are often unneeded for business activities, but provide an avenue of attack and compromise of the organization.	13
Attackers exploit wireless access points to gain entry into a target organization's internal network, as well as exploit wireless client systems to steal sensitive information.	14
Attackers who gain access to internal enterprise systems gather and exfiltrate sensitive information without detection by the victim organization.	15
Attackers exploit poorly designed network architectures by locating unneeded or unprotected connections, weak filtering, or a lack of separation of important systems or business functions.	16
Attackers compromise target organizations that do not exercise their defenses to determine and continually improve their effectiveness.	17

Attackers operate undiscovered in organizations without effective incident response capabilities, and when they are discovered, such organizations often cannot properly contain the attack, eradicate the attackers' presence, and recover to a secure production state.	18
Attackers who compromise systems may alter important data, potentially jeopardizing organizational effectiveness via polluted information.	19
Attackers exploit users and system administrators via social engineering scams that work because of a lack of security skills and awareness.	20