

Twenty Critical Security Controls for Effective Cyber Defense: Consensus Audit Guidelines (CAG)

Version 3.1 October 3, 2011

Table of Contents

- Twenty Critical Security Controls for Effective Cyber Defense: Consensus Audit Guidelines (CAG)..... 1
- Version 3.1 October 3, 2011* 1
- Introduction** 3
 - Forging Consensus among Chief Information Security Officers, Chief Information Officers, and Inspectors General on Technical Requirements for System Administrators and Security Personnel.. 5
 - Relationship of the 20 Critical Controls to National Institute of Standards and Technology Guidelines 6
 - Relationship of the 20 Critical Controls to the Australian Government’s Defence Signals Directorate 35 Strategies to Mitigate Targeted Cyber Intrusions 7
 - Relationship of the 20 Critical Controls to the National Security Agency’s Associated Manageable Network Plan Revision 2.1 Milestones and Network Security Tasks..... 8
 - Document Contributors..... 8
- The 20 Critical Controls**..... 9
 - Insider versus Outsider Threats 10
 - Relationship to Other US Federal Guidelines, Recommendations, and Requirements..... 12
 - Periodic and Continual Testing of Controls 12
 - Future Evolution of the 20 Critical Controls 12
- Description of Controls** 13
 - Critical Control 1: Inventory of Authorized and Unauthorized Devices..... 13
 - Critical Control 2: Inventory of Authorized and Unauthorized Software..... 16
 - Critical Control 3: Secure Configurations for Hardware and Software on Laptops, Workstations, and Servers 19
 - Critical Control 4: Continuous Vulnerability Assessment and Remediation 23
 - Critical Control 5: Malware Defenses..... 26
 - Critical Control 6: Application Software Security 29
 - Critical Control 7: Wireless Device Control..... 33
 - Critical Control 8: Data Recovery Capability..... 36
 - Critical Control 9: Security Skills Assessment and Appropriate Training to Fill Gaps 37
 - Critical Control 10: Secure Configurations for Network Devices such as Firewalls, Routers, and Switches..... 39
 - Critical Control 11: Limitation and Control of Network Ports, Protocols, and Services..... 42
 - Critical Control 12: Controlled Use of Administrative Privileges 44

Critical Control 13: Boundary Defense	48
Critical Control 14: Maintenance, Monitoring, and Analysis of Audit Logs	52
Critical Control 15: Controlled Access Based on the Need to Know	55
Critical Control 16: Account Monitoring and Control	57
Critical Control 17: Data Loss Prevention	60
Critical Control 18: Incident Response Capability	63
Critical Control 19: Secure Network Engineering.....	65
Critical Control 20: Penetration Tests and Red Team Exercises	66
Summary and Action Plan.....	69
Appendix A: Mapping between the 20 Critical Security Controls and National Institute of Standards and Technology Special Publication 800-53, Revision 3, Priority 1 Items.....	70
Appendix B: Mapping between the 20 Critical Security Controls and Australian Government’s DSD 35 Mitigation Strategies	72
Appendix C: Attack Types	75

Introduction

Securing the United States against cyber attacks has become one of the nation's highest priorities. To achieve this objective, networks and systems, as well as the operations teams that support them, must vigorously defend against a variety of internal and external threats. To respond to those attacks that *are* successful, defenses must be prepared to detect and thwart follow-on attacks on internal enterprise networks as attackers spread inside a compromised network. A critical component of such a defense system is continuous monitoring—that is, the ability to automatically test and validate whether current security measures are working and proactively remediate vulnerabilities in a timely manner.

A central tenet of cyber defense is that “offense must inform defense.” In other words, knowledge of actual attacks that have compromised systems provides the essential foundation upon which to build effective defenses. The US Senate Homeland Security and Government Affairs Committee made this a central tenet of the U.S. ICE Act of 2009 (the proposed new Federal Information Security Management Act –FISMA). The legislation calls on federal agencies to “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.” The legislation further calls for the White House to ensure that these steps are taken.

Because federal agencies do not have unlimited resources, current and past federal chief information officers (CIOs) and chief information security officers (CISOs) have concluded that the only rational way to meet these requirements is to jointly establish a **prioritized baseline of information security measures and controls** that can be continuously monitored using automated mechanisms.

As we look to the future, it is clear that ongoing initiatives within the federal government will continue to expand interconnectivity across organizations to better support citizens and internal government operations. Interconnectivity means that security vulnerabilities in one area of a particular federal agency can become the path to compromise other parts of the federal system. This exposure also exists for commercial organizations that have connections across divisions and with partners and other networks. An exposure in one network can be used to tunnel to other connected networks, making it easier to compromise a network from behind the firewall. Cloud computing also increases an organization's footprint and exposure. It is essential that a prioritized set of overarching security controls be established that can be applied to the enterprise environments in which organizations operate, and across the federal government system as well.

This consensus document of 20 Critical Controls aims to begin the process of establishing a prioritized baseline of information security measures and controls that can be applied across federal and commercial environments. The consensual effort that produced this document identifies 20 specific technical security controls that are viewed as effective in blocking currently known high-priority attacks as well as those attack types expected in the near future. Each of the 20 control areas includes multiple individual subcontrols that specify actions an organization can take to help improve its defenses.

The control areas and individual subcontrols focus on various technical aspects of information security with the primary goal of helping organizations prioritize their efforts to defend 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 many other 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 nontechnical, aspects of information security. Organizations should build a comprehensive approach to 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 the following:

- Defenses should focus on addressing the most common and damaging attack activities occurring today, and on those anticipated in the near future.
- Enterprise environments must ensure that consistent controls are in place across the organization to effectively negate attacks.
- Defenses should be automated where possible and periodically or continuously measured using automated measurement techniques where feasible.
- A variety of specific technical activities should be undertaken to produce a more consistent defense against attacks that occur on a frequent basis against numerous organizations.
- Root cause problems must be fixed in order to ensure the prevention or timely detection of attacks
- Metrics should be established that facilitate common ground for measuring the effectiveness of security measures, providing a common language for executives, information technology specialists, auditors, and security officials to communicate about risk within the organization.

The controls presented here are also designed to support organizations with different levels of information security capabilities. To help organizations design a sound security baseline and then improve beyond that baseline, subcontrols included in each of the summaries of the 20 Critical Controls have been grouped into specific categories:

- *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 these subcontrols do not necessarily provide comprehensive protection against the most critical attacks. The intent of identifying “quick wins” is to highlight where security can be improved rapidly.
- *Improved visibility and attribution.* These subcontrols focus on improving the process, architecture, and technical capabilities of organizations so that they can monitor their networks and computer systems and better visualize their own IT operations. Attribution is associated with determining which computer systems, and potentially which users, are generating specific events. Such improved visibility and attribution helps organizations detect attack attempts, locate the points of entry for successful attacks, identify already-compromised machines, interrupt infiltrated attackers' activities, and gain information

about the sources of an attack. In other words, these controls improve an organization's situational awareness of its environment. These subcontrols are identified in this document as "Visibility/Attribution."

- *Hardened configuration and improved information security hygiene.* These subcontrols are designed to improve an organization's information security stance by reducing the number and magnitude of potential security vulnerabilities and by improving the operations of networked computer systems. They focus 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. These subcontrols are identified in this document as "Configuration/Hygiene."
- *Advanced.* These subcontrols are designed to further improve the security of an organization beyond the other three categories. Organizations already following all of the other subcontrols should focus on this category.

In general, organizations should compare all 20 control areas against their current status and develop an organization-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 quick wins through visibility/attribution and configuration/hygiene and up to advanced. As a start, organizations with limited information security programs may want to address the quick wins subcontrols in order to make rapid progress and build momentum within their information security program.

Many of these controls can be implemented and measured using existing tools found in many government agencies and corporations. Other controls can be implemented 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 objectives associated with the most important elements of the given control. Each control area also includes a test section that provides information about how organizations can evaluate their implementation of each control metric. These tests are devised to support automation wherever possible so that organizations can achieve reliable, scalable, and continuous measurements of their adherence to the controls and related metrics.

Forging Consensus among Chief Information Security Officers, Chief Information Officers, and Inspectors General on Technical Requirements for System Administrators and Security Personnel

CISOs and CIOs are charged with improving the state of information security across the federal government and commercial organizations. Moreover, they are spending increasing amounts of money to secure their systems. However, securing those systems is enormously complex, and so there is a need to focus attention and resources on the most critical risk (and therefore the highest payoff) areas.

CISOs and CIOs also want and need specific guidance that can be consistently applied enterprise-wide and upon which their performance in improving security can be consistently and fairly evaluated. At the same time, federal inspectors general (IGs), lawyers, executives, and auditors want to ensure that CIOs and CISOs are doing what is necessary to secure systems, but these authorities also need specific guidance on how to measure security. Finally, technical personnel associated with information security operations and system administration require a specific set of technical activities to help them defend 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 provide to their technical system administration and information security personnel to ensure that their 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 successfully launched against federal systems, industrial-base systems, and commercial networks.

This effort also draws on the success and insights from developing and using standardized concepts to identify, communicate, and document security-relevant characteristics/data. These standards include the common identification of vulnerabilities and their severity, definition of secure configurations, inventory of systems and platforms, and identification of application weaknesses. The standards have emerged over the last decade through collaborative research and deliberation among 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), which was sponsored by the National Institute of Standards and Technology (NIST) and mandated for the Federal Desktop Core Configuration (FDCC). SCAP uses 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 National Institute of Standards and Technology Guidelines

Revision 3 of NIST Special Publication 800-53 (*Recommended Security Controls for Federal Information Systems*) provides a comprehensive set of excellent security controls. By contrast, the document presented here 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, the controls do map directly to about one-third of the 145 controls identified in NIST Special Publication 800-53 (see Appendix A). In fact, the mapping shows that the 20 Critical Controls are a proper subset of the Priority 1 items in 800-53. A mapping for each of the 20 Critical Controls to the specific set of 800-53 controls is

provided in the text below and a complete mapping is included as Appendix A of this document. Moreover, the attack-based analysis of the 20 Critical Controls confirms that they are the most critical subset of the NIST Special Publication 800-53 control catalog. Once organizations have addressed the 20 Critical Controls, it is recommended that they use 800-53 to ensure that they have assessed and implemented an appropriate set of management controls as well as additional technical controls that address specific risk areas.

The first step for CIOs and CISOs who seek to improve cyber security is to assess the 20 Critical Controls as a baseline set of “common controls” for their organizations as defined by 800-53. The basis for this recommendation is that the consensus process used to develop the 20 Critical Controls, as well as pilot efforts by the US State Department, have validated that the controls correlate to the highest technical and operational threat areas for the 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 an organization’s “high water mark” when assessing the enterprise-wide potential security risk of confidentiality, integrity, or availability of interconnected systems and information within the organization’s enterprise environment. Once the CIO and CISO have forged consensus, it is recommended that the 20 Critical Controls be adopted as the foundation for an organization’s technical and operational controls. Similarly, the 20 Critical Controls would also serve as a primary basis for future security audits and evaluations. If an organization’s CIO and CISO determine that their environment warrants 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 unlikely that an organization with Internet connectivity would determine that the 20 Critical Controls are not applicable to its operations. However, if an organization’s CIO and CISO determine that some of the controls are not applicable, this also can be documented using processes outlined in 800-53.

Relationship of the 20 Critical Controls to the Australian Government’s Defence Signals Directorate 35 Strategies to Mitigate Targeted Cyber Intrusions

In 2010, the Australian Defence Signals Directorate (DSD) developed a list of 35 prioritized mitigation strategies to defend networks and systems from cyber attack. In 2011, the DSD updated and reprioritized this list based on a detailed analysis of recent cyber attacks across the Australian Government. Based on this analysis, the DSD estimates that these defenses, implemented together, would have prevented at least 85% of the intrusions into the networks they analyzed.

The DSD’s 35 Mitigation Strategies focus on individual tasks organizations can undertake to improve their security stance, and, as such, they map directly to the Top 20 Critical Controls. In fact, the 35 Mitigation Strategies are a focused subset of the Top 20 Critical Controls, as shown in detail in Appendix B.

Relationship of the 20 Critical Controls to the National Security Agency's Associated Manageable Network Plan Revision 2.1 Milestones and Network Security Tasks

The Associated Manageable Network Plan Revision 2.1 Milestones and Network Security Tasks, developed by the National Security Agency (NSA), is a series of steps designed to improve an unmanageable and insecure network. The plan is intended to be a long-term solution, as implementing the milestones may take a significant amount of resources and time (possibly months or even years). A key component of the 20 Critical Controls is to ensure that the security of a network can be maintained in a cost-effective manner. Therefore there is a close mapping and correlation between the 20 Critical Controls and the Manageable Network Plan.

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. First-hand knowledge and input on how computer and network attacks are being carried out and the defensive techniques that are most important to thwart them was provided by a wide range of people and organizations, including:

1. Blue team members inside the Department of Defense (DoD) 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 that identify prior intrusions while conducting vulnerability assessment activities.
3. US Computer Emergency Readiness Team staff and other nonmilitary 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 have been compromised.
4. Military investigators who fight cyber crime.
5. The FBI and other law enforcement 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 to determine how the attackers penetrated the systems and what they did subsequently.
8. Red team members inside the 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 more than 100 other collaborators has been incorporated into the current version of this document. To assemble the 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 entitled "How do attackers exploit the absence of this control?" Furthermore, Appendix C of this

document provides a list of each of the attack types that fueled the development of the 20 Critical Controls.

The 20 Critical Controls

The 20 Critical Controls were agreed upon by knowledgeable individuals from the above-mentioned groups. The list includes controls that can be continuously monitored and validated at least in part in an automated manner and some that must be validated manually. It is important to note that the 20 control categories are prioritized based on the NSA attack mitigation scores. The process of gathering these specific controls and subcontrols focused on identifying the highest priority defenses and represents a subset of controls found in other audit guidelines and documents. Each of the 20 control areas is important and offers high-priority techniques for thwarting real-world attacks.

Critical Controls subject to 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. Continuous Vulnerability Assessment and Remediation
5. Malware Defenses
6. Application Software Security
7. Wireless Device Control
8. Data Recovery Capability (validated manually)
9. Security Skills Assessment and Appropriate Training to Fill Gaps (validated manually)
10. Secure Configurations for Network Devices such as Firewalls, Routers, and Switches
11. Limitation and Control of Network Ports, Protocols, and Services
12. Controlled Use of Administrative Privileges
13. Boundary Defense
14. Maintenance, Monitoring, and Analysis of Security Audit Logs
15. Controlled Access Based on the Need to Know
16. Account Monitoring and Control
17. Data Loss Prevention
18. Incident Response Capability (validated manually)
19. Secure Network Engineering (validated manually)
20. Penetration Tests and Red Team Exercises (validated manually)

The pages that follow more fully describe each of these controls.

The descriptions include how attackers currently exploit the absence of the control; a list of detailed subcontrols that outline what an organization needs to do in each area and the requirements for measuring these activities; and suggestions regarding how standardized measurements can be applied. After organizations implement the controls and gain experience with automation, the document can be used as an audit guide that CIOs can use to ensure that

they are taking the right actions for effective cyber defense, and that IGs can use to verify the CIOs' tests.

Insider versus Outsider Threats

A quick review of the 20 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 the 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 and well-funded actors such as nation-states may sometimes employ attack techniques that require extreme defenses that go beyond the scope of this document.

The controls are not limited to blocking only the initial compromise of systems, but also to detecting already-compromised machines and preventing or disrupting an 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 superuser 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 back doors and root kits. Attackers with access to machines can also cause damage by stealing, altering, or destroying information; impairing the system's functionality in order to jeopardize its business effectiveness or mission; or using the compromised machine as a jump-off point to compromise 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 Figure 1 are covered by the controls described in this document. Defenses in any of the rings help limit the abilities of attackers, but improved defenses are required across all three rings and their intersections. It is important to note that the 20 Critical Controls are designed to help improve defenses across each of these rings, rather than to merely prevent initial compromise.

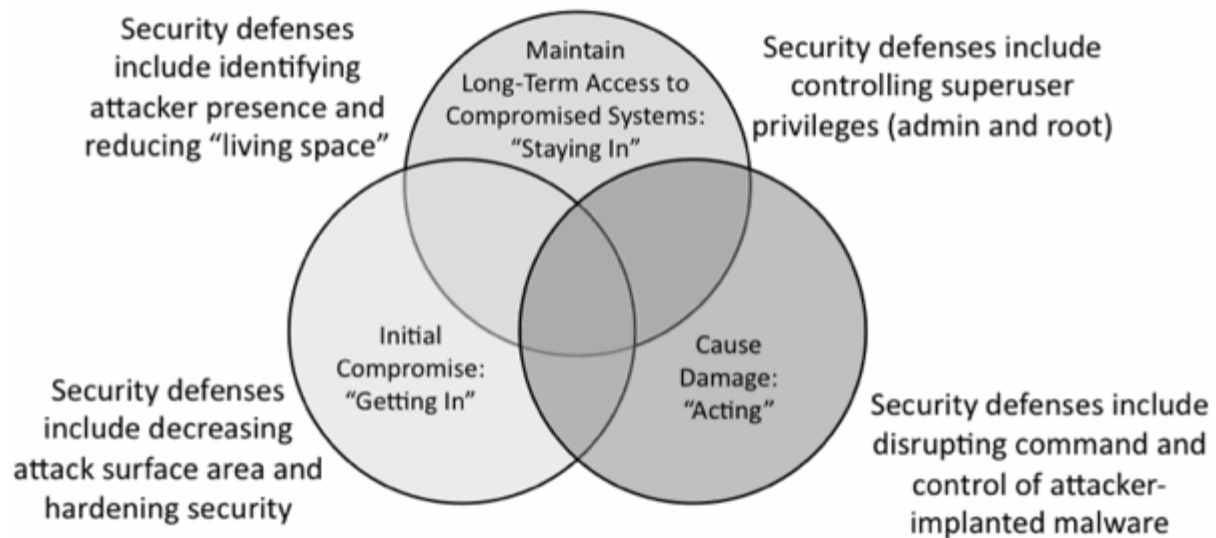


Figure 1: Computer Attacker Activities and Associated Defenses

Attack patterns are also constantly changing, and it is important that organizations be fully aware of their environments and that they have automated controls in place to provide continuous monitoring. One area of increased focus for attackers is compromises against critical infrastructure (i.e., domain name systems and border gateway protocols). As more countries become involved with maintaining the Internet and Internet Protocol version 6 (IPv6), these problems will increase in severity.

Zero day attacks are also increasing. While it is difficult to protect against zero day attacks directly, in many cases there is extraneous software running that is not required. In the Aurora attack—named for the file path on the attacker’s machine that was included in two of the malware binaries believed to be associated with the attack—client-side exploitation occurred through a vulnerability in older software. Computer attackers used this mechanism to compromise Google and numerous other large organizations in an attempt to access sensitive information on their internal networks. While Aurora was a zero-day exploit, there is likely little business need to have Internet Explorer 6 installed in an environment.

Advanced persistent threats using targeted attacks are also increasing. In the Stuxnet attack on industrial software and equipment, a targeted worm was capable of causing significant harm. Many people think worms are a thing of the past, but stealthy, covert worms are still in use because they provide an automated mechanism for attackers. Worms can be highly targeted and extremely malicious, emphasizing the need for a variety of controls.

Organizations typically focus their security on systems that are visible from the Internet or other untrusted networks. However those systems are often connected to internal systems that potentially contain sensitive information. Attackers are frequently using pivot points, bouncing from system to system until they achieve their goal, be it operational impact or information theft.

Relationship to Other US Federal Guidelines, Recommendations, and Requirements

The 20 Critical 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, SCAP, FDCC, FISMA, manageable network plans, 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 Special Publication 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. Appendix A maps the individual controls in this document to specific recommendations of NIST Special Publication 800-53.

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 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 easily and consistently shared among organizations 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 test stages, 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 work. Without appropriate authorization, background checks, and possibly clearance, such tests may be impossible. Such tests should also be supervised or reviewed by appropriate organization 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 20 Critical Controls

The consensus effort to define critical security controls is an evolving process. 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, replicable, measurable, scalable, and reliable process throughout the federal government.

Description of Controls

Critical Control 1: Inventory of Authorized and Unauthorized Devices

How Do Attackers Exploit the Absence of this Control?

Many criminal groups and nation-states deploy systems that continuously scan address spaces of target organizations, waiting for new and 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 accessible via the Internet. 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.

As new technology continues to come out, many employees bring personal devices into work and connect them to the network. These devices could already be compromised and be used to infect internal resources. Attackers are also increasing the use of pivot points, compromising one system and using that as an anchor point to break into other systems that might not be directly visible to the attacker.

How to Implement, Automate, and Measure the Effectiveness of this Control

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. Quick wins: 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. Visibility/Attribution: 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 Internet protocol (IP) address on the network, including, but not

limited to desktops, laptops, servers, network equipment (routers, switches, firewalls, etc.), printers, storage area networks, voiceover-IP telephones, etc.

3. Visibility/Attribution: The asset inventory created must also include data on whether the device is a portable device. Devices such as mobile phones, tablets, laptops, and other portable electronic devices that store or process data must be identified, regardless of whether or not they are attached to the organization's network.
4. Visibility/Attribution: 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.
5. Configuration/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.
6. Configuration/Hygiene: In addition to an inventory of hardware, organizations should develop an inventory of information assets that 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.
7. Configuration/Hygiene: Deploy network level authentication via 802.1x to limit and control which devices can be connected to the network. 802.1x must be tied into the inventory data to determine authorized vs. unauthorized systems.
8. Advanced: Network access control can be used to monitor authorized systems so if attacks occur, the impact can be remediated by moving the untrusted system to a virtual local area network that has minimal access.

Associated NIST Special Publication 800-53, Revision 3, Priority 1 Controls

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

Associated NSA Manageable Network Plan Milestones and Network Security Tasks

Milestone 2: Map Your Network

Milestone 3: Network Architecture

Personal Electronic Device (PED) Management

Procedures and Tools to Implement and Automate 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 media access control (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. Once MAC addresses are confirmed, switches should implement 802.1x to only allow authorized systems to connect to the network.

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 transmission control protocol (TCP) synchronize (SYN) or acknowledge (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 virtual private network (VPN) technology may appear on the network for a time, and then be disconnected from it. Whether physical or virtual, each machine 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 of alerting or sending e-mail 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 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 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 10 locations on the network, including a selection of subnets associated with demilitarized zones (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.

Control 1 Sensors, Measurement, and Scoring

Sensor: Automated asset inventory system

Measurement: Look for tools such as Sourcefire Network RNA, GFI Network Inventory Management Tool to have been deployed and operating.

Score: Score is based on how frequently and recently scans are being and have been performed.

Sensor: Network-level authentication

Measurement: Verify that 802.1x or similar proprietary solution has been deployed to manage asset connectivity. Solutions such as Cisco Identity Based Networking.

Score: Score is the percentage of ports in the enterprise that are managed.

Critical Control 2: Inventory of Authorized and Unauthorized Software

How Do Attackers Exploit the Absence 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 a system is compromised. 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 to Implement, Automate, and Measure the Effectiveness of this Control

1. Quick wins: 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. Visibility/Attribution: 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.
3. Visibility/Attribution: The software inventory 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.
4. Configuration/Hygiene: Deploy application white listing technology that allows systems to run only approved software and prevents execution of all other software on the system, based on an automatically generated list of valid software from a representative sample machine. Such white listing tools must be based on acceptable hashing algorithms for determining authorized binaries to execute on a system.
5. Advanced: Virtual machines and/or air-gapped systems should also be used to isolate and run applications that are required but based on higher risk and that should not be installed within a networked environment.
6. Advanced: Configure client workstations with non-persistent virtualized operating environments that can be quickly and easily restored to a trusted snapshot on a periodic basis.

Associated NIST Special Publication 800-53, Revision 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

Associated NSA Manageable Network Plan Milestones and Network Security Tasks

Milestone 7: Baseline Management
Executable Content Restrictions

Procedures and Tools to Implement and Automate 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 white and black lists of programs allowed to run or blocked from executing are included in many modern endpoint security suites. Moreover, commercial solutions are increasingly bundling together anti-virus, anti-spyware, personal firewall, and host-based intrusion detection systems (IDS) and intrusion prevention systems (IPS), along with application white 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 white and black lists based on executable path, hash, or regular expression matching. Some even include a gray-list function that allows administrators to define rules for execution of specific programs only by certain users and at certain times of day, and black lists 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 e-mail. 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 state of security, in the future organizations should strive for even more rapid alerting and isolation, with notification about unauthorized software 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 10 systems

on the network. Two of the systems must be included in the asset inventory database, while the other systems do not need to be included. 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.

Control 2 Sensors, Measurement, and Scoring

Sensor: Software inventory system

Measurement: Scan systems on a monthly basis and determine the number of unauthorized pieces of software that are installed. Verify that if an unauthorized piece of software is found one month, it is removed from the system the next.

Score: 100 percent if no unauthorized software is found. Minus 1 percent for each piece of unauthorized software that is found. If the unauthorized software is not removed, minus 2 percent each consecutive month.

Sensor: Application white listing software

Measurement: Run application white listing on all key servers and review the logs once a month. Determine the number of expectations that are made or the number of servers it is disabled on.

Score: Pass if there are less than 25 exceptions a month and less than 15 systems that have the software turned off. Otherwise fail.

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

How Do Attackers Exploit the Absence 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 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 extraneous services that are 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 to Implement, Automate, and Measure the Effectiveness of this Control

1. Quick wins: Strict configuration management should be followed, building a secure image that is used to build all new systems that are deployed to the enterprise. Any existing system that becomes compromised is re-imaged with the secure build. Regular updates to this image are integrated into the organization's change management processes.
2. Quick wins: System images must have documented security settings that are tested before deployment, approved by an organization change control board, and registered with a central image library for the organization or multiple organizations. These images should be validated and refreshed on a regular basis (e.g., every six months) to update their security configuration in light of recent vulnerabilities and attack vectors.
3. Quick wins: Standardized images should represent hardened versions of the underlying operating system and the applications installed on the system, such as those released by the NIST, NSA, Defense Information Systems Agency (DISA), Center for Internet Security (CIS), and others. This hardening would typically include removal of unnecessary accounts, disabling or removal of unnecessary services, and configuring non-executable stacks and heaps through the use of operating system features such as Data Execution Prevention (DEP). 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 erecting host-based firewalls.
4. Quick wins: Any deviations from the standard build or updates to the standard build should be documented and approved in a change management system.
5. Quick wins: Organizations 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.
6. Quick wins: The master images themselves must be stored on securely configured servers, with integrity checking tools and change management to ensure that only authorized changes to the images are possible. Alternatively, these master images can be stored in offline 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.
7. Quick wins: Run the last version of software and make sure it is fully patched. Remove outdated or older software from the system.
8. Configuration/Hygiene: All remote administration of servers, workstation, network devices, and similar equipment shall be done over secure channels. Protocols such as telnet, VNC, RDP, or other protocols that do not natively support strong encryption should only be used if they are performed over a secondary encryption channel, such as SSL or IPSEC.
9. Configuration/Hygiene: At least once a month, run assessment programs on a varying sample of systems to determine which ones are configured according to the secure configuration guidelines.

10. Configuration/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.
11. Configuration/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.
12. Configuration/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 Special Publication 800-53, Revision 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

Associated NSA Manageable Network Plan Milestones and Network Security Tasks

Milestone 7: Baseline Management
Configuration and Change Management

Procedures and Tools to Implement and Automate this Control

Organizations can implement this control by developing a series of images and secure storage servers for hosting these standard images. Commercial and/or free configuration management tools can then be employed to measure the settings operating system and applications of managed machines 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 log in 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 e-mail notification to a list of enterprise administrative personnel. Systems must block installation, prevent execution, or quarantine unauthorized software within one additional hour, alerting or sending e-mail when this action has occurred. Every 24 hours after that point, the system must alert or send e-mail about the status of the system until it has been removed from the network or remediated. While the 24-hour and one-hour timeframes represent the current metric to help organizations improve their state of security, in the future organizations should strive for even more rapid alerting and isolation, with notification about unauthorized changes sent within two minutes 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, but that does contain additional services, ports and configuration files changes, onto the network. This must be performed on 10 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 e-mail 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 e-mail notification process.

Control 3 Sensors, Measurement, and Scoring

Sensor: File integrity software

Measurement: File integrity monitoring software is deployed on servers as a part of the base configuration. Centralized solutions like Tripwire are preferred over stand-alone solutions.

Score: 50 percent awarded for using a solution like Tripwire with a central monitoring/reporting component. The remaining 50 percent is based on the percentage of servers on which the solution is deployed.

Sensor: Standard images

Measurement: Standard images for the installation of systems have been created based on an accepted security standard published by organizations such as CIS, NSA, DISA and others.

Score: Pass/Fail

Sensor: Network-based image deployment system

Measurement: Computers are built from secured masters pushed out by image servers. Solutions such as Acronis, Ghost, and others are appropriate.

Score: Percentage of systems built from and potentially managed by the solution.

Critical Control 4: Continuous Vulnerability Assessment and Remediation

How Do Attackers Exploit the Absence 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 dangerous 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 to Implement, Automate, and Measure the Effectiveness of this Control

1. Quick wins: 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. Any vulnerability identified should be remediated in a timely manner, with critical vulnerabilities fixed within 48 hours.
2. Quick wins: 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. Second, 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
3. Visibility/Attribution: Organizations should deploy automated patch management tools and software update tools for operating system and third-party software on all systems for which such tools are available and safe.
4. Configuration/Hygiene: In order to overcome limitations of unauthenticated vulnerability scanning, organizations should ensure that all vulnerability scanning is performed in authenticated mode 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.
5. Configuration/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 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.

6. Configuration/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 that minimum standards are met.
7. Configuration/Hygiene: Security personnel should chart the numbers of unmitigated, critical vulnerabilities for each department/division.
8. Configuration/Hygiene: Security personnel should share vulnerability reports indicating critical issues with senior management to provide effective incentives for mitigation.
9. Configuration/Hygiene: Organizations should measure the delay in patching new vulnerabilities and ensure that the delay is equal to or less than the benchmarks set forth by the organization.
10. Configuration/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.

Associated NIST Special Publication 800-53, Revision 3, Priority 1 Controls

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

Associated NSA Manageable Network Plan Milestones and Network Security Tasks

Milestone 6: Patch Management

Procedures and Tools to Implement and Automate 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 organization or even across organizations, 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 log in 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, organizations 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 inadvertent introduction of security weaknesses by administrators.

Effective organizations link their vulnerability scanners with problem-ticketing systems that automatically monitor and report progress on fixing problems, and that make unmitigated critical vulnerabilities visible 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 elapses 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 by the vendor or administrator. Because of this, all patch checks should reconcile system patches with a list of patches each vendor has announced on its website.

Control 4 Metric

All machines identified by the asset inventory system associated with Critical 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 e-mail 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 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 sent within two minutes.

Control 4 Test

To evaluate the implementation of Control 4 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 finish.

Control 4 Sensors, Measurement, and Scoring

Sensor: Vulnerability scanner

Measurement: Tools such as Tenable's Security Center, Qualysguard, Secunia and others should be deployed and configured to run automatically.

Score: Pass or fail.

Critical Control 5: Malware Defenses

How Do Attackers Exploit the Absence of this Control?

Malicious software is an integral and dangerous aspect of Internet threats, targeting end-users and organizations via web browsing, e-mail 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 to Implement, Automate, and Measure the Effectiveness of this Control

1. Quick wins: Organizations should employ automated tools to continuously monitor workstations, servers, and mobile devices for active, up-to-date anti-malware protection with anti-virus, anti-spyware, personal firewalls, and host-based IPS functionality. All malware detection events should be sent to enterprise anti-malware administration tools and event log servers.
2. Quick wins: 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. Quick wins: 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 serial advanced technology attachment devices, mounted network shares, or other removable media.
4. Quick wins: Organizations should configure systems so that they conduct an automated anti-malware scan of removable media when it is inserted.
5. Quick wins: All attachments entering the organization’s e-mail gateway should be scanned and blocked if they contain malicious code or file types unneeded for the organization’s business. This scanning should be done before the e-mail is placed in the user’s inbox. This includes email content filtering and web content filtering.
6. Visibility/Attribution: Automated monitoring tools should use behavior-based anomaly detection to complement and enhance traditional signature-based detection.
7. Configuration/Hygiene: Advanced: Organizations should deploy network access control tools to verify security configuration and patch-level compliance before granting access to a network.
8. Advanced: Continuous monitoring should be performed on outbound traffic. Any large transfers of data or unauthorized encrypted traffic should be flagged and, if validated as malicious, the computer should be moved to an isolated VLAN.
9. Advanced: Organizations should implement an Incident Response process which allows their IT Support Organization to supply their Security Team with samples of malware running undetected on corporate systems. Samples should be provided to the Anti-Virus vendor for ‘out-of-band’ signature creation and deployed to the enterprise by system administrators.

Associated NIST Special Publication 800-53, Revision 3, Priority 1 Controls

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

Associated NSA Manageable Network Plan Milestones and Network Security Tasks

Virus Scanners and Host Intrusion Prevention Systems (HIPS)
Personal Electronic Device (PED) Management
Network Access Protection/Control (NAP/NAC)
Security Gateways, Proxies, and Firewalls
Network Security Monitoring

Procedures and Tools to Implement and Automate 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 capable of consistently handling this task. 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 in-

depth security, 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 follow-on investigation.

Control 5 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 e-mail 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 state of security, in the future organizations should strive for even more rapid detection and malware isolation, with notification about malware in the enterprise sent within two minutes and blocking, execution prevention, or quarantine actions taken within five minutes.

Control 5 Test

To evaluate the implementation of Control 5 on a periodic basis, the evaluation team must move a benign software test program that appears to be malware (such as an EICAR file or benign hacker tools) but that is not included in the official authorized software list to 10 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 indicating that the software has been blocked or quarantined is received within one hour. 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 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 this time transferring the benign malware to 10 systems via e-mail instead. The organization must expect the same notification results as noted with the removable media test.

Control 5 Sensors, Measurement, and Scoring

Sensor: Anti-virus management. TrendMicro, Symantec, McAfee and Kaspersky all have management consoles that can validate configurations and run reports.

Measurement: (1) Determine if anti-virus program is running on all systems; (2) Confirm that it is configured to run whenever a file is opened or attempted to run on the system.

Score: Determine the percent of systems that are running anti-virus programs and the percent of systems that are properly configured and average the two together.

Sensor: Honeypots deployed. While there are some programs that can perform system emulation (i.e., Honeyd), virtual machines are usually utilized to create honeypots.

Measurement: Number of connections to the honeypot correlated to the number of unique IP addresses the connection is coming from.

Score: None in a five-day period: 100 percent; 10 in a five-day period: 90 percent; 20 in a five-day period: 80 percent.

Sensor: Patch management software. Microsoft WSUS can be used but only works with Microsoft products. BigFix, Lumension, Shavlik, and LANDesk can also be used.

Measurement: Confirm that every system in the asset inventory database is running patch management software. Validate that all systems are receiving/applying patches within two weeks after a patch is released. For critical patches it should be within 48 hours.

Score: 100 percent if all systems are running patch management software and fully patched. Minus 1 percent for each system not running patch management software, and 2 percent for each system that is not receiving patches in a timely manner.

Critical Control 6: Application Software Security

How Do Attackers Exploit the Absence 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, cross-site scripting, cross-site request forgery, and click-jacking of code to gain control over vulnerable machines. In one attack, 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 it.

How to Implement, Automate, and Measure the Effectiveness of this Control

1. Quick wins: 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, specific application firewalls should be deployed if such tools are available for the given application type. If the traffic is encrypted, the device should either sit behind the encryption or be capable of decrypting the traffic prior to analysis. If neither option is appropriate, a host-based web application firewall should be deployed.
2. Visibility/Attribution: At a minimum, explicit error checking should be done for all input. Whenever a variable is created in source code, the size and type should be determined. When input is provided by the user it should be verified that it does not exceed the size or the data type of the memory location in which it is stored or moved in the future.
3. Configuration/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.
4. Configuration/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.
5. Configuration/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.
6. Configuration/Hygiene: Organizations should verify that security considerations are taken into account throughout the requirements, design, implementation, testing, and other phases of the software development life cycle of all applications.
7. Configuration/Hygiene: Organizations should ensure that all software development personnel receive training in writing secure code for their specific development environment.
8. Configuration/Hygiene: Require that all in-house-developed software include white-list filtering capabilities for all data input and output associated with the system. These white lists 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.
9. Configuration/Hygiene: Sample scripts, libraries, components, compilers, or any other unnecessary code that is not being used by an application should be uninstalled or removed from the system.

Associated NIST Special Publication 800-53, Revision 3, Priority 1 Controls

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

Associated NSA Manageable Network Plan Milestones and Network Security Tasks

Milestone 3: Network Architecture

Milestone 7: Baseline Management

Security Gateways, Proxies, and Firewalls

Procedures and Tools to Implement and Automate 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 and application penetration testing expertise. The Common Weakness Enumeration (CWE) initiative is used 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” published free online by Mitre and the SANS Institute 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, Mitre’s Common Attack Pattern Enumeration and Classification can be used to organize and record the breadth of the testing for the CWEs and to enable testers to think like attackers in their development of test cases.

Control 6 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 implementing a compensating control) within 15 days of discovery of the flaw.

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

Control 6 Test

To evaluate the implementation of Control 6 on a monthly basis, an evaluation team must use a web application vulnerability scanner to test for each type of flaw identified in the regularly updated list of the “25 Most Dangerous Programming Errors” by Mitre and the SANS Institute. 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 that 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 what happened. 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 finish.

Control 6 Sensors, Measurement, and Scoring

Sensor: Web Application Firewall

Measurement: Verify that a web application firewall (WAF) is installed between applications and users. Products such as F5 Application Security Manager, ModSecurity, Art of Defence Hyperguard, and Trustwave WebDefend are recommended.

Score: Automated tool/process verifies: WAF is installed and functioning: 50 points. WAF configuration covers OWASP top 10: 20 points. WAF configuration defends against top 25 programming errors: 30 points.

Sensor: Web application firewall

Measurement: Central logging tool shows evidence that logs are being collected from WAF.

Score: Automated tool/process periodically verifies that WAF is generating logs into the security event manager or similar: 100 points. Failure to identify log entries = 0.

Sensor: Vulnerability/Configuration testing tools are running and reporting automatically. (Critical Control 10)

Measurement: Configuration and targets for vulnerability management tools used to satisfy Critical Control 10 appropriately configured to monitor application and application-base OS configuration issues.

Score: Automated tool verifies that configuration and target list for Critical Control 10 includes application servers: 100 points. Failure to identify application server = 0.

Critical Control 7: Wireless Device Control

How Do Attackers Exploit the Absence of this Control?

Major thefts of data 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 cafes. 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 to Implement, Automate, and Measure the Effectiveness of this Control

1. Quick wins: 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. Quick wins: 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. Quick wins: 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. Visibility/Attribution: Organizations should use wireless intrusion detection systems (WIDS) to identify rogue wireless devices and detect attack attempts and successful compromises. In addition to WIDS, all wireless traffic should be monitored by a wired IDS as traffic passes into the wired network.
5. Visibility/Attribution: 802.1x should be used to control which devices are allowed to connect to the wireless network.
6. Visibility/Attribution: A site survey should be performed to determine what areas within the organization need coverage. After the wireless access points are strategically placed, the signal strength should be tuned to minimize leakage to areas that do not need coverage.
7. Configuration/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.

8. Configuration/Hygiene: For devices that do not have an essential wireless business purpose, organizations should disable wireless access in the hardware configuration (basic input/output system or extensible firmware interface), with password protections to lower the possibility that the user will override such configurations.
9. Configuration/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.
10. Configuration/Hygiene: Organizations should ensure that all wireless traffic leverages at least Advanced Encryption Standard (AES) encryption used with at least WiFi Protected Access 2 protection.
11. Configuration/Hygiene: Organizations should ensure that wireless networks use authentication protocols such as Extensible Authentication Protocol-Transport Layer Security (EAP/TLS) or Protected Extensible Authentication Protocol (PEAP), which provide credential protection and mutual authentication.
12. Configuration/Hygiene: Organizations should ensure that wireless clients use strong, multi-factor authentication credentials to mitigate the risk of unauthorized access from compromised credentials.
13. Configuration/Hygiene: Organizations should disable peer-to-peer wireless network capabilities on wireless clients, unless such functionality meets a documented business need.
14. Configuration/Hygiene: Organizations should disable wireless peripheral access of devices (such as Bluetooth), unless such access is required for a documented business need.
15. Configuration/Hygiene: Wireless access points should never be directly connected to the private network. They should either be placed behind a firewall or put on a separate VLAN so all traffic can be examined and filtered.
16. 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 Special Publication 800-53, Revision 3, Priority 1 Controls

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

Associated NSA Manageable Network Plan Milestones and Network Security Tasks

Remote Access Security

Procedures and Tools to Implement and Automate this Control

Effective organizations run commercial wireless scanning, detection, and discovery tools as well as commercial wireless intrusion detection systems.

Additionally, the security team should 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 organization network.

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

Control 7 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 e-mail 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 state of security, in the future organizations should strive for even more rapid alerting and isolation, with notification about an unauthorized wireless devices sent within two minutes and isolation within five minutes.

Control 7 Test

To evaluate the implementation of Control 7 on a periodic basis, the evaluation team staff must configure 10 unauthorized but hardened wireless clients and wireless access points to the organization's network and attempt to connect them to its 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 service set identifier 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 any of the above-noted systems attempt 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.

Control 7 Sensors, Measurement, and Scoring

Sensor: Wireless access point

Measurement: Determine if any rogue access points are connected to the network.

Score: 100 percent if no rogue access points are detected for two months. Minus 5 percent for each unauthorized access point that is discovered.

Sensor: Wireless IDS

Measurement: Utilizing the asset inventory database, determine if any clients are trying to make a connection to an access point that they are not authorized to make.

Score: 100 percent if there are no unauthorized connections attempted by clients. Minus 2 percent for each unauthorized client connection.

Sensor: Wireless vulnerability scanner

Measurement: Perform scans of all wireless access points on a monthly basis looking for known vulnerabilities or unauthorized configuration changes.

Score: 100 percent if no unauthorized changes are found. Minus 2 percent for known vulnerabilities and minus 5 percent for known vulnerabilities that were previously fixed or for unauthorized configuration changes.

Critical Control 8: Data Recovery Capability

How Do Attackers Exploit the Absence 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 are discovered, it can be extremely difficult for organizations without a trustworthy data recovery capability to remove all aspects of the attacker's presence on the machine.

How to Implement, Automate, and Measure the Effectiveness of this Control

1. Quick wins: 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, the operating system, application software, and data on a machine should each be included in the overall backup procedure. These three components of a system do not have to be included in the same backup file or use the same backup software. However, each must be backed up at least weekly.

2. Quick wins: Data on backup media should be tested on a regular basis by performing a data restoration process to ensure that the backup is properly working.
3. Quick wins: Key personnel should be trained on both the backup and restoration processes. To be ready in case a major incident occurs, alternative personnel should also be trained on the restoration process just in case the primary IT point of contact is not available.
4. Configuration/Hygiene: Organizations should ensure that backups are properly protected via physical security or encryption when they are stored locally, as well as when they are moved across the network.
5. Configuration/Hygiene: Backup media, such as hard drives and tapes, should be stored in physically secure, locked facilities.

Associated NIST Special Publication 800-53, Revision 3, Priority 1 Controls

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

Associated NSA Manageable Network Plan Milestones and Network Security Tasks

Backup Strategy

Procedures and Tools to Implement and Automate 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.

Control 8 Sensors, Measurement, and Scoring

Sensor: Backup software

Measurement: Verify that an automated backup solution is in place for all critical systems. The automated system could be a tape library system, a hot-spare network file store, or something similar.

Score: Percentage of critical systems that are backed up. The score diminishes based on the number of days since the last successful backup of a critical system.

Critical Control 9: Security Skills Assessment and Appropriate Training to Fill Gaps

Skills of Five Groups of People Constantly Being Tested 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 alteration would have on their mission.

Any organization that hopes to be ready to find and respond to attacks effectively owes it to its employees and contractors to find the gaps in its knowledge and 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.

Training is also closely tied to policy and awareness. Policies tell people what to do, training provides them the skills to do it, and awareness changes behaviors so that people follow the policy. Training should be mapped against the skills required to perform a given job. If after training, users are still not following the policy, that policy should be augmented with awareness.

How to Implement, Automate, and Measure the Effectiveness of this Control

1. Quick wins: 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, and should present proven defenses against the latest attack techniques.
2. Quick wins: Awareness should be carefully validated with policies and training. Policies tell users what to do, training provides them the skills to do it, and awareness changes their behavior so that they understand the importance of following the policy.
3. Visibility/Attribution: Metrics should be created for all policies and measured on a regular basis. Awareness should focus on the areas that are receiving the lowest compliance score.
4. Configuration/Hygiene: Organizations should devise periodic security awareness assessment quizzes to be given to employees and contractors on at least an annual basis in order to determine whether they understand the information security policies and procedures, as well as their role in those procedures.
5. Configuration/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.
6. Advanced: Provide awareness sessions for users who are not following policies after they have received appropriate training.

Associated NIST Special Publication 800-53, Revision 3, Priority 1 Controls

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

Associated NSA Manageable Network Plan Milestones and Network Security Tasks

Training

Procedures and Tools to Implement and Automate this Control

The key to upgrading skills is measurement—not through certification examinations, but through 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 to improve their skills. In addition, the organization can develop training programs that directly fill the gaps and maintain employee readiness.

Control 9 Sensors, Measurement, and Scoring

Sensor: Policy

Measurement: For each policy statement, measure the overall compliance every month.

Score: Pass if the compliance for each policy statement is increasing, fail if it is decreasing for any statement for more than three months in a row.

Critical Control 10: Secure Configurations for Network Devices such as Firewalls, Routers, and Switches

How Do Attackers Exploit the Absence 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, as the exceptions are deployed, and as 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 neither properly analyzed nor 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 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 to Implement, Automate, and Measure the Effectiveness of this Control

1. Quick wins: 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 organization 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. Quick wins: At network interconnection points—such as Internet gateways, inter-organization connections, and internal network segments with different security controls—implement ingress and egress filtering to allow only those ports and protocols with an explicit and documented business need. All other ports and protocols should be blocked with default-deny rules by firewalls, network-based IPS, and/or routers.
3. Quick wins: Network devices that filter unneeded services or block attacks (including firewalls, network-based IPS, 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. Configuration/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 IPS, 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. Configuration/Hygiene: Network filtering technologies employed between networks with different security levels (firewalls, network-based IPS tools, and routers with access controls lists) should be deployed with capabilities to filter Internet Protocol version 6 (IPv6) traffic. However, if IPv6 is not currently being used it should be disabled. Since many operating systems today ship with IPv6 support activated, filtering technologies need to take it into account.
6. Configuration/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. Configuration/Hygiene: The latest stable version of a network device’s inter-network operating system (IOS) or firmware must be installed within 30 days of the update being released from the device vendor.
8. 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, on entirely different physical connectivity for management sessions for network devices.

Associated NIST Special Publication 800-53, Revision 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

Associated NSA Manageable Network Plan Milestones and Network Security Tasks

Milestone 7: Baseline Management
Configuration and Change Management

Procedures and Tools to Implement and Automate 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 access controls lists (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 10 Metric

The system must be capable of identifying any changes to network devices, including routers, switches, firewalls, and 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 e-mail 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 10 Test

To evaluate the implementation of Control 10 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 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 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 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.

Control 10 Sensors, Measurement, and Scoring

Sensor: File integrity software

Measurement: File integrity monitoring software is deployed on all network devices or run across the network as a part of the base configuration. Centralized solutions like Tripwire are preferred over stand-alone solutions.

Score: 50 percent awarded for using a solution like Tripwire with a central monitoring/reporting component. The remaining 50 percent is based on the percentage of servers on which the solution is deployed.

Sensor: Standard images

Measurement: Standard images for the installation of systems have been created based on an accepted security standard published by organizations such as CIS, NSA, DISA, and others.

Score: Pass/Fail

Sensor: Packet generation tools

Measurement: Confirm that the network infrastructure properly handles, routes and filters IPv6 traffic.

Score: Pass or Fail.

Critical Control 11: Limitation and Control of Network Ports, Protocols, and Services

How Do Attackers Exploit the Absence 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 domain name system (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 to Implement, Automate, and Measure the Effectiveness of this Control

1. Quick wins: 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. Quick wins: Automated port scans should be performed on a regular basis against all key servers and compared to a known effective baseline. If a new port is found open, an alert should be generated and reviewed.
3. Visibility/Attribution: Any server that is visible from the Internet or an untrusted network should be verified, and if it is not required for business purposes it should be moved to an internal VLAN and given a private address.
4. Configuration/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. Sometimes services are turned on for projects or limited engagements, and should be turned off when they are no longer needed.
5. Configuration/Hygiene: Operate critical services on separate physical host machines, such as DNS, file, mail, web, and database servers.
6. Advanced: Application firewalls should be placed in front of any critical servers to verify and validate the traffic going to the server. Any unauthorized services or traffic should be blocked and an alert generated.

Associated NIST Special Publication 800-53, Revision 3, Priority 1 Controls

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

Associated NSA Manageable Network Plan Milestones and Network Security Tasks

Milestone 3: Network Architecture
Security Gateways, Proxies, and Firewalls

Procedures and Tools to Implement and Automate 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 11 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 e-mail 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 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 state of security, in the future organizations should strive for even more rapid alerting, with notification about an unauthorized open port on the network sent within two minutes.

Control 11 Test

To evaluate the implementation of Control 11 on a periodic basis, the evaluation team must install hardened test services with network listeners on 10 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 team must verify that the system provides details of the location of all of the systems where test services have been installed.

Control 11 Sensors, Measurement, and Scoring

Sensor: Host-based firewalls

Measurement: Verify that all systems have a host-based firewall installed and operating.

Score: Percentage of systems for which the host based firewall can be verified to be functioning.

Sensor: Automated network scans

Measurement: Ensure that the solutions for Critical Controls 3 and 10 are being leveraged to monitor changes to services on protected systems.

Score: Pass/Fail.

Critical Control 12: Controlled Use of Administrative Privileges

How Do Attackers Exploit the Absence of this Control?

The misuse of administrator privileges is a primary 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, running as a privileged user, is fooled into opening a malicious e-mail attachment, downloading and opening a file from a malicious website, 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. Similar attacks occur with e-mail. An administrator inadvertently opens an e-mail that contains an infected attachment and this is used to obtain a pivot point within the network that is used to attack other systems.

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 to Implement, Automate, and Measure the Effectiveness of this Control

1. Quick wins: Organizations should use automated tools to inventory all administrative accounts and validate that each person with administrative privileges on desktops, laptops, and servers is authorized by a senior executive and that all administrative passwords have at least 12 pseudorandom characters, consistent with the FDCC standard.
2. Quick wins: 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. Quick wins: Organizations should configure all administrative-level accounts to require regular password changes on a frequent interval of no longer than 60 days.
4. Quick wins: 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, at a frequent interval of no longer than 90 days.
5. Quick wins: Passwords for all systems should be stored in a well-hashed or encrypted format, with weaker formats such as Windows LANMAN hashes eliminated from the environment. Furthermore, files containing these encrypted or hashed passwords required for systems to authenticate users should be readable only with superuser privileges.
6. Quick wins: Organizations should use automated scripts to ensure that administrator accounts are used only for system administration activities, and not for reading e-mail, composing documents, or surfing the Internet. Web browsers and e-mail clients especially must be configured to never run as administrator.
7. Quick wins: Through policy and user awareness, organizations should require that administrators establish unique, different passwords for their administrator and non-administrative accounts. Each person requiring administrative access should be given his/her own separate account. Administrative accounts should never be shared. Users should only use the Windows “administrator” or Unix “root” accounts in emergency situations. Domain administration accounts should be used when required for system administration instead of local administrator accounts.
8. Quick wins: Organizations should configure operating systems so that passwords cannot be re-used within a certain timeframe, such as six months.
9. Visibility/Attribution: Organizations should implement focused auditing on the use of administrative privileged functions and monitor for anomalous behavior (e.g., system reconfigurations during the night shift).
10. Visibility/Attribution: 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. Configuration/Hygiene: All administrative access, including domain administrative access, should use two-factor authentication.

12. Configuration/Hygiene: Access to a machine (either remotely or locally) should be blocked for administrator-level accounts. Instead, administrators should be required to access a system using a fully logged and non-administrative account. Then, once logged in to the machine without administrative 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. Each user would use their own administrator account and enter a password each time that is different than their user account.
13. Configuration/Hygiene: If services are outsourced to third parties, language should be included in the contracts to ensure that they properly protect and control administrative access. It should be validated that they are not sharing passwords and have accountability to hold administrators liable for their actions.
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 Special Publication 800-53, Revision 3, Priority 1 Controls

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

Associated NSA Manageable Network Plan Milestones and Network Security Tasks

Milestone 5: User Access

Milestone 7: Baseline Management

Procedures and Tools to Implement and Automate 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 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 that 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 12 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 superuser 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 state of security, in the future organizations should strive for even more rapid alerting, with notification about new additions to superuser groups sent within two minutes.

Control 12 Test

To evaluate the implementation of Control 12 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 10 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 superuser 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 10 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.

Control 12 Sensors, Measurement, and Scoring

Sensor: Password assessment tool

Measurement: Automated password assessment is performed and reports generated.

L0phtcrack 6 provides the ability to schedule periodic password assessments and generate reports. Other tools can be scripted to provide similar capabilities. Systems from Critical Control 10 must also be leveraged to check for default passwords on all networked systems.

Score: Automatically verify that password assessment tool is installed: 20 points.

Automatically verify that password assessment tool configured or scripted to run automatically: 40 points. Vulnerability scan configuration from Critical Control 10 automatically verified to cover default passwords on all networked systems: 40 points.

Sensor: SEIM

Measurement: Log management and reporting systems from Critical Control 6 must be collecting information on the use of administrative credentials. This requires the ability to report on this criterion, and the individual systems must be configured to report on the use of administrative credentials.

Score: Automated tool verifies systems are configured to report the use of credentials with administrative privileges. 100 points.

Critical Control 13: Boundary Defense

How Do Attackers Exploit the Absence of this Control?

Attackers focus on exploiting systems that they can reach across the Internet, including 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 police content by looking for attacks and evidence of compromised machines, boundary defenses should be multi-layered, relying on firewalls, proxies, DMZ perimeter networks, and network-based IPS and IDS. It is also critical to filter both inbound and outbound traffic.

It should be noted that boundary lines between internal and external networks are diminishing as a result of 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 of boundaries, effective security deployments still rely on carefully configured boundary defenses that separate networks with different threat levels, sets of users, and levels of control. Even with the blurring of internal and external networks, effective multi-layered defenses of perimeter networks help lower the number of successful attacks, allowing security personnel to focus on attackers who have devised methods to bypass boundary restrictions.

How to Implement, Automate, and Measure the Effectiveness of this Control

The boundary defenses included in this control build on Critical Control 10. The additional recommendations here focus 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 an intruder's access to the other parts of the network.

1. Quick wins: Organizations should deny communications with (or limit data flow to) known malicious IP addresses (black lists) or limit access to trusted sites (white lists). Tests can be periodically carried out by sending packets from bogon source IP addresses 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. Quick wins: Deploy network-based IDS sensors on Internet and extranet DMZ systems and networks that look for unusual attack mechanisms and detect compromise of these systems. These network-based IDS sensors may detect attacks through the use of signatures, network behavior analysis, or other mechanisms to analyze traffic.
3. Quick wins: Network-based IPS devices should be deployed to compliment IDS by blocking known bad signature or behavior of attacks. As attacks become automated, methods such as IDS typically delay the amount of time it takes for someone to react to an attack. A properly configured network-based IPS can provide automation to block bad traffic.
4. Quick wins: 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. This traffic should be sent to a properly configured Security Event Information Management (SEIM) system so that events can be correlated from all devices on the network.
5. Quick wins: To lower the chance of spoofed e-mail messages, implement the Sender Policy Framework (SPF) by deploying SPF records in DNS and enabling receiver-side verification in mail servers.
6. Visibility/Attribution: Define a network architecture that clearly separates internal systems from DMZ and extranet systems. DMZ systems are machines that need to communicate with the internal network as well as the Internet, while extranet systems are those whose primary communication is with other systems at a business partner. DMZ systems should never contain sensitive data and internal systems should never be directly accessible from the Internet.
7. Visibility/Attribution: Design and implement network perimeters so that all outgoing web, file transfer protocol (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 black list; and applying white lists of allowed sites that can be accessed through the proxy while blocking all other sites. Organizations should force outbound traffic to the Internet through an authenticated proxy server on the enterprise perimeter. Proxies can also be used to encrypt all traffic leaving an organization.
8. Visibility/Attribution: 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.

9. Configuration/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.
10. Configuration/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.
11. Configuration/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.
12. Configuration/Hygiene: Organizations should develop plans to rapidly deploy filters on internal networks to help stop the spread of malware or an intruder.
13. Advanced: To minimize the impact of an attacker pivoting between compromised systems, only allow DMZ systems to communicate with private network systems via application proxies or application-aware firewalls over approved channels
14. 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 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 sessions.

Associated NIST Special Publication 800-53, Revision 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

Associated NSA Manageable Network Plan Milestones and Network Security Tasks

Milestone 3: Network Architecture
Security Gateways, Proxies, and Firewalls
Remote Access Security
Network Security Monitoring

Procedures and Tools to Implement and Automate this Control

One element of this control can be implemented using free or commercial IDS 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 Hypertext Transfer Protocol (HTTP) traffic that bypasses HTTP proxies. By sampling traffic regularly, such as over a three-hour period once per week, information security personnel can 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 can 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 that packets take from the sender to the receiver system.

Control 13 Metric

The system must be capable of identifying any unauthorized packets sent into or out of a trusted zone and ensure that the packets are properly blocked and/or trigger alerts. Any unauthorized packets must be detected within 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 state of security, in the future organizations should strive for even more rapid alerting, with notification about unauthorized packets in a trusted zone sent within two minutes.

Control 13 Test

To evaluate the implementation of Control 13 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 indicating that the unauthorized traffic is now being blocked is received within one hour. 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 state where it does not forward traffic when it crashes or becomes flooded.

Control 13 Sensors, Measurement, and Scoring

Sensor: Network-based IDS

Measurement: Verify that a network-based IDS has been deployed at all network boundaries. Sourcefire, Sguil, Palo Alto Networks IPS, etc. are appropriate solutions.

Score: Pass/Fail

Sensor: Network-based IPS

Measurement: Verify that a network-based IPS has been deployed to supplement the network-based IDS. Systems such as Sourcefire are appropriate solutions.

Score: Pass/Fail.

Sensor: Public/Screened network packet logs

Measurement: Verify that a packet logging solution has been deployed at the perimeter for all public/screened networks. Solutions include daemonlogger, tcpdump, IDSBench and similar systems.

Score: Number of days of logs / 90 * 100, not to exceed 100 percent.

Sensor: Proxy Servers

Measurement: Proxy servers have been deployed to the screened network to act as an intermediary between trusted and untrusted systems.

Score: Pass/Fail.

Critical Control 14: Maintenance, Monitoring, and Analysis of Audit Logs

How Do Attackers Exploit the Absence 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 have been compromised, without protected and complete logging records they are blind to the details of the attack and to 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 nonexistent 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 to Implement, Automate, and Measure the Effectiveness of this Control

1. Quick wins: 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 initiative. If systems cannot generate logs in a standardized

format, log normalization tools can be deployed to convert logs into a standardized format.

2. Quick wins: 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. The logs must be archived and digitally signed on a periodic basis.
3. Quick wins: All remote access to a network, whether to the DMZ or the internal network (i.e., VPN, dial-up, or other mechanism), should be logged verbosely.
4. Quick wins: 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. Failed logon attempts must also be logged.
5. Quick wins: Security personnel and/or system administrators should run biweekly reports that identify anomalies in logs. They should then actively review the anomalies, documenting their findings.
6. Visibility/Attribution: Each organization should include at least two synchronized time sources (i.e., Network Time Protocol – NTP) from which all servers and network equipment retrieve time information on a regular basis so that timestamps in logs are consistent.
7. Visibility/Attribution: Network boundary devices, including firewalls, network-based IPS, and inbound and outbound proxies, should be configured to verbosely log all traffic (both allowed and blocked) arriving at the device.
8. Visibility/Attribution: For all servers, organizations should ensure that 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.
9. Visibility/Attribution: Organizations should deploy a SEIM system tool for log aggregation and consolidation from multiple machines and for log correlation and analysis. Standard government scripts for analysis of the logs should be deployed and monitored, and customized local scripts should also be used. Using the SEIM tool, 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.

Associated NIST Special Publication 800-53, Revision 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)

Associated NSA Manageable Network Plan Milestones and Network Security Tasks

Remote Access Security
Log Management

Procedures and Tools to Implement and Automate 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 in the event a follow-up investigation is 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 in order to ensure that each managed item actively connected to the network is periodically generating logs.

Analytical programs such as SEIM for reviewing logs can be useful, but the capabilities employed to analyze audit logs are quite extensive, including just a cursory examination by a person. Actual correlation tools can make audit logs far more useful for subsequent manual inspection. 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 14 Metric

The system must be capable of logging all events across the network. The logging must be validated across both network-based 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 state of security, in the future organizations should strive for even more rapid alerting, with notification about a logging failure sent within two minutes.

Control 14 Test

To evaluate the implementation of Control 14 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, 10 servers, and 10 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.

Control 14 Sensors, Measurement, and Scoring

Sensor: NTP

Measurement: Confirm that NTP is being used to synchronize time for all devices and that all clocks are in synch.

Score: Pass or fail.

Sensor: Vulnerability scanner

Measurement: Run a vulnerability scanner against random servers utilizing nonintrusive scans. Determine whether the information appeared in the logs.

Score: Pass or fail.

Sensor: SEIM

Measurement: Correlate logs to a central source and determine that all servers are properly logging.

Score: 100 percent if all systems are properly logging. Minus 5 percent for each system that is not logging.

Critical Control 15: Controlled Access Based on the Need to Know

How Do Attackers Exploit the Absence 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 to Implement, Automate, and Measure the Effectiveness of this Control

1. Quick wins: Organizations should establish a multi-level data identification/classification scheme (e.g., a three- or four-tiered scheme with data separated into categories based on the impact of exposure of the data).
2. Quick wins: Organizations should ensure that file shares have defined controls (such as Windows share access control lists) that specify at least that only “authenticated users” can access the share.
3. Visibility/Attribution: Organizations should enforce detailed audit logging for access to nonpublic data and special authentication for sensitive data.
4. Configuration/Hygiene: The network should be segmented based on the trust levels of the information stored on the servers. Whenever information flows over a network of lower trust level, the information should be encrypted.

5. Configuration/Hygiene: The use of portable USB drives should either be limited or data should automatically be encrypted before it is written to a portable drive.
6. Advanced: Host-based data loss prevention (DLP) should be used to enforce ACLs even when data is copied off a server. In most organizations, access to the data is controlled by ACLs that are implemented on the server. Once the data have been copied to a desktop system the ACLs are no longer enforced and the users can send the data to whomever they want.
7. Advanced: Deploy honeytokens on key servers to identify users who might be trying to access information that they should not access.

Associated NIST Special Publication 800-53, Revision 3, Priority 1 Controls

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

Associated NSA Manageable Network Plan Milestones and Network Security Tasks

Milestone 3: Network Architecture

Milestone 5: User Access

Procedures and Tools to Implement and Automate this Control

It is important that an organization understand what its sensitive information is, where it resides, and who needs access to it. To derive sensitivity levels, organizations need to put together a list of the key types of data and the overall importance to the organization. This analysis would be used to create an overall data classification scheme for the organization. At a base level, a data classification scheme is broken down into two levels: public (unclassified) and private (classified). Once the private information has been identified, it can then be further subdivided based on the impact it would have to the organization if it was compromised.

Once the sensitivity of the data has been identified, it needs to be traced back to business applications and the physical servers that house those applications. The network then needs to be segmented so that systems of the same sensitivity level are on the same network and segmented from systems of different trust levels. If possible, firewalls need to control access to each segment. If data are flowing over a network of a lower trust level, encryption should be used.

Job requirements should be created for each user group to determine what information the group needs access to in order to perform its jobs. Based on the requirements, access should only be given to the segments or servers that are needed for each job function. Detailed logging should be turned on for all servers so that access can be tracked and that situations where someone is accessing data that they should not be accessing can be examined.

Control 15 Metric

The system must be capable of detecting all attempts by users to access files on local systems or network-accessible file shares without the appropriate privileges, and it 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 state of security, in the future organizations should strive for even more rapid alerting, with notification about unauthorized access attempts sent within two minutes.

Control 15 Test

To evaluate the implementation of Control 15 on a periodic basis, the evaluation team must create two test accounts each on 10 representative systems in the enterprise: five server machines and five client systems. For each system evaluated, 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 nonprivileged 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.

Control 15 Sensors, Measurement, and Scoring

Sensor: ACLs

Measurement: Verify that ACLs are properly configured on all file shares. Additionally, verify that the ACLs restrict access to groups with the appropriate need to know. Tools such as ShareEnum or SoftPerfect's NetworkScanner can be used to identify shares and extract ACLs automatically.

Score: Automated tool scans for file shares and reports ACLs inappropriately configured. Anonymous or public shares within a controlled environment result in a zero score: 100 points.

Sensor: DLP

Measurement: Tools such as McAfee Host DLP are deployed on all systems containing or having access to controlled data.

Score: Automated tool verifies that DLP software is installed and functioning.

Sensor: Honeytokens

Measurement: Honeytokens are powerful for the early detection of intentional and accidental data leakage. Honeytokens should be deployed on all sensitive data stores. The identity of the honeytokens must be kept closely guarded.

Score: Automated tool periodically verifies the existence of honeytokens on all file shares.

Sensor: IDS/DLP with honeytokens

Measurement: To be effective, it is necessary to monitor for the transmission of honeytokens.

Score: Automated tool verifies that IDS and DLP configurations are capable of detecting transmission or access to honeytokens.

Critical Control 16: Account Monitoring and Control

How Do Attackers Exploit the Absence 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 to Implement, Automate, and Measure the Effectiveness of this Control

1. Quick wins: Review all system accounts and disable any account that cannot be associated with a business process and owner.
2. Quick wins: 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. Quick wins: Organizations should establish and follow a process for revoking system access by disabling accounts immediately upon termination of an employee or contractor.
4. Quick wins: Organizations should regularly monitor the use of all accounts, automatically logging off users after a standard period of inactivity.
5. Quick wins: 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. Quick wins: 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.
7. Quick wins: All non-administrator accounts should be required to have a minimum length of 12 characters, contain letters, numbers, and special characters, be changed at least every 90 days, have a minimal age of one day, and not be allowed to use the previous 15 passwords as a new password.
8. Quick wins: After eight failed logon attempts within a 45-minute period, the account should be locked for 120 minutes.
9. Visibility/Attribution: 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.
10. Visibility/Attribution: Organizations should monitor attempts to access deactivated accounts through audit logging.
11. Configuration/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 percent. This includes flagging the use of user's credentials from a computer other than computers usually used by the user.

Associated NIST Special Publication 800-53, Revision 3, Priority 1 Controls

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

Associated NSA Manageable Network Plan Milestones and Network Security Tasks

Milestone 5: User Access

Procedures and Tools to Implement and Automate 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 use home-grown scripts or third-party log analysis tools to analyze this information and profile user access of various systems.

Accounts must also be tracked very closely. Any account that is dormant must be disabled and eventually removed from the system. All active accounts must be traced back to authorized users of the system and it must be ensured that their passwords are robust and changed on a regular basis. Users must also be logged out of the system after a period of no activity to minimize the possibility of an attacker using their system to extract information from the organization.

Control 16 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 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 sent within two minutes.

Control 16 Test

To evaluate the implementation of Control 16 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 30 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.

Control 16 Sensors, Measurement, and Scoring

Sensor: Account management software

Measurement: Using management tools like Microsoft System Center, Trusted Computer Solutions Security Blanket, Intellitactics Security Manager, Qwest Enterprise Security Reporter, and MaxPowerSoft AD Reports, determine the number of current accounts on the system.

Perform validation that all active accounts are valid and baseline the list. Create a process so that whenever a new account is authorized and added or removed, the system baseline list is also updated. Once a week, compare the list of current accounts to the baseline list and flag any anomalies that exist.

Score: 100 percent if there are no unauthorized accounts created for a six-month period. Minus 1 percent for each unauthorized account that exists.

Sensor: Account management software

Measurement: Using management tools like Microsoft System Center, Trusted Computer Solutions Security Blanket, Intellitactics Security Manager, Qwest Enterprise Security Reporter, and MaxPowerSoft AD Reports, scan all active accounts and flag any accounts that have a default password or have not been logged into for 60 days and are still active.

Score: 100 percent if there are no active accounts that should be disabled. Minus 1 percent for each active account that should be disabled.

Sensor: Account management software

Measurement: Using management tools like Microsoft System Center, Trusted Computer Solutions Security Blanket, Intellitactics Security Manager, Qwest Enterprise Security Reporter, and MaxPowerSoft AD Reports, determine the number of failed logon attempts.

Score: 100 percent if there are no failed logon attempts. Minus 1 percent for each failed logon attempt.

Critical Control 17: Data Loss Prevention

How Do Attackers Exploit the Absence 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 and a potential threat to national security. While some data are 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 nonexistent.

The phrase “data loss prevention” 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, and 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 to Implement, Automate, and Measure the Effectiveness of this Control

1. Quick wins: Organizations should deploy approved hard drive encryption software to mobile machines that hold sensitive data.
2. Visibility/Attribution: 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. Visibility/Attribution: Deploy an automated tool on network perimeters that monitors for certain sensitive information (i.e., personally identifiable information), keywords, and other document characteristics to discover unauthorized attempts to exfiltrate data across network boundaries and block such transfers while alerting information security personnel.
4. Visibility/Attribution: Conduct periodic scans of server machines using automated tools to determine whether sensitive data (i.e., personally identity, health, credit card, and classified information) is present on the system in clear text. These tools, which search for patterns that indicate the presence of sensitive information, can help identify if a business or technical process is leaving behind or otherwise leaking sensitive information in data at rest.
5. Visibility/Attribution: Use outbound proxies to be able to monitor and control all information leaving an organization.
6. Configuration/Hygiene: Data should be moved between networks using secure, authenticated, and encrypted mechanisms.
7. Configuration/Hygiene: Data stored on removable and 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 are automatically encrypted without user intervention.
8. Configuration/Hygiene: 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, enterprise software should be used that can configure systems to allow only specific USB devices (based on serial number or other unique property) to be accessed, and that can automatically encrypt all data placed on such devices. An inventory of all authorized devices must be maintained.

9. Configuration/Hygiene: Use network-based DLP solutions to monitor and control the flow of data within the network. Any anomalies that exceed the normal traffic patterns should be noted and appropriate action taken to address them.
10. Advanced: Monitor all traffic leaving the organization and detect any unauthorized use of encryption. Attackers often use an encrypted channel to bypass network security devices. Therefore it is essential that organizations are able to detect rogue connections, terminate the connection, and remediate the infected system.

Associated NIST Special Publication 800-53, Revision 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

Associated NSA Manageable Network Plan Milestones and Network Security Tasks

Personal Electronic Device (PED) Management

Data-at-Rest Protection

Network Security Monitoring

Procedures and Tools to Implement and Automate 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 17 Metric

The system must be capable of identifying unauthorized data leaving the organization, whether via network file transfers or removable media. Within one hour of a data exfiltration event or attempt, 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, the system owner must also be included 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 state of security, in the future organizations should strive for even more rapid alerting, with notification about data exfiltration events or attempts sent within two minutes.

Control 17 Test

To evaluate the implementation of Control 17 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 removable media. Each of the following tests must be performed at least three times:

- Attempt to transfer large data sets across network boundaries from an internal system.
- Attempt to transfer test data sets of personally identifiable information (that trigger DLP systems but do not contain sensitive data) across network boundaries from an internal system (using multiple keywords specific to the business).
- Attempt to maintain a persistent network connection for at least 10 hours across network boundaries between an internal and external system, even though little data may be exchanged.
- Attempt to maintain a network connection across network boundaries using an anomalous service port number between an internal and external system.
- Insert a USB token into an organization system and attempt 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.

Control 17 Sensors, Measurement, and Scoring

Sensor: Network-based DLP tool

Measurement: Verify that a reputable DLP solution has been installed and configured on the network.

Score: Pass/Fail

Sensor: Data encryption

Measurement: Verify that a full disk encryption solution has been deployed for all mobile systems that handle sensitive data.

Score: Percentage of mobile systems with full disk encryption installed.

Sensor: Behavioral network-based IDS

Measurement: Ensure that the solution used for Critical Control 5 is capable of performing behavioral analysis to identify unusual outbound data flows.

Score: Pass/Fail.

Critical Control 18: Incident Response Capability

How Do Attackers Exploit the Absence of this Control?

Considerable 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 plans 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 a far greater impact, causing more damage, infecting more systems, and possibly exfiltrating more sensitive data than would otherwise be possible were an effective incident response plan in place.

NIST Special Publication 800-61 contains detailed guidelines for creating and running an incident response team (available at <http://csrc.nist.gov/publications/nistpubs/800-61-rev1/SP800-61rev1.pdf>).

How to Implement, Automate, and Measure the Effectiveness of this Control

1. Quick wins: Organizations should ensure that they have written incident response procedures that 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. Quick wins: Organizations should assign job titles and duties for handling computer and network incidents to specific individuals.
3. Quick wins: Organizations should define management personnel who will support the incident handling process by acting in key decision-making roles.
4. Quick wins: Organizations should devise organization-wide standards for the time required for system administrators and other personnel to report anomalous events to the incident handling team, the mechanisms for such reporting, and the kind of information that should be included in the incident notification. This reporting should also include notifying the appropriate US Community Emergency Response Team in accordance with all government requirements for involving that organization in computer incidents.
5. Quick wins: Organizations should publish information for all personnel, including employees and contractors, regarding reporting computer anomalies and incidents to the incident handling team. Such information should be included in routine employee awareness activities.
6. Configuration/Hygiene: Organizations should conduct periodic incident scenario sessions for personnel associated with the incident handling team to ensure that they understand current threats and risks, as well as their responsibilities in supporting the incident handling team.

Associated NIST Special Publication 800-53, Revision 3, Priority 1 Controls

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

Associated NSA Manageable Network Plan Milestones and Network Security Tasks

Incident Response and Disaster Recovery Plans
Training

Procedures and Tools to Implement and Automate 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.

Control 18 Sensors, Measurement, and Scoring

Sensor: Incident response plan

Measurement: Simulate an incident and determine how quickly the team responds and remediates the issue.

Score: Compare the actual results with the expected results and take the overall percent.

Critical Control 19: Secure Network Engineering

How Do Attackers Exploit the Absence 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 to Implement, Automate, and Measure the Effectiveness of this Control

1. Quick wins: The network should be designed using a minimum of a three-tier architecture (DMZ, middleware, and private network). Any system accessible from the Internet should be on the DMZ, but DMZ systems never contain sensitive data. Any system with sensitive data should reside on the private network and never be directly accessible from the Internet. DMZ systems should communicate with private network systems through an application proxy residing on the middleware tier.
2. Configuration/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. Visibility/Attribution: 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. Visibility/Attribution: Security should be built into all phases of the software development lifecycle, ensuring that any security issues are addressed as early as possible.

5. Configuration/Hygiene: 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 Special Publication 800-53, Revision 3, Priority 1 Controls

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

Associated NSA Manageable Network Plan Milestones and Network Security Tasks

Milestone 3: Network Architecture

Procedures and Tools to Implement and Automate this Control

To help ensure a consistent, defensible network, the architecture of each network should be based on a template that describes the network's overall layout and the services it provides. Organizations should prepare 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.

Control 19 Sensors, Measurement, and Scoring

Sensor: Port or vulnerability scanner

Measurement: Determine which systems are visible from the Internet or untrusted systems. Sensitive systems or databases should not be accessible from untrusted networks.

Score: 100 percent if no unauthorized systems are found. Minus 10 percent for unauthorized systems and minus 15 percent for database servers or any system that contains sensitive information.

Sensor: Network diagram

Measurement: Check and scan the network to determine that it matches the network diagram.

Score: Minus 5 percent for each unauthorized change.

Critical Control 20: Penetration Tests and Red Team Exercises

How Do Attackers Exploit the Absence 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 a vulnerability assessment. 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.

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 and objective insights about the existence of vulnerabilities and about the efficacy of defenses and mitigating controls already in place and even those planned for future implementation.

How to Implement, Automate, and Measure the Effectiveness of this Control

1. Quick wins: Organizations should conduct regular external and internal 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 as from within its boundaries (i.e., on the internal network) to simulate both outsider and insider attacks.
2. Visibility/Attribution: 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. Visibility/Attribution: Organizations should ensure that systemic problems discovered in penetration tests and red team exercises are fully mitigated.
4. Visibility/Attribution: Organizations should measure how well the organization has reduced the significant enablers for attackers by setting up automated processes to find:
 - Cleartext e-mails 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 finding 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. Visibility/Attribution: Social engineering should be included within a penetration test. The human element is often the weakest link in an organization and one that attackers often target.
6. Advanced: Organizations should devise a scoring method for determining the results of red team exercises so that results can be compared over time.
7. 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 supervisory control and data acquisition and other control systems.

Associated NIST Special Publication 800-53, Revision 3, Priority 1 Controls

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

Procedures and Tools to Implement and Automate 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 a minimum, 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-value 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.

Control 20 Sensors, Measurement, and Scoring

Sensor: Automated penetration testing tool

Measurement: Determine the number of systems that have vulnerabilities that can be exploited and determine to what level they can be exploited.

Score: 100 percent if no vulnerabilities can be exploited. Minus 5 percent for guest-level exploitation, 10 percent for user-level access, and 15 percent for root-level access.

Summary and Action Plan

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 itself 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.

Your Action Plan

Given that these critical controls so closely track current threats and attacks, we recommend that CIOs and CISOs consider several immediate actions to ensure the effectiveness of their security programs:

- 1) Conduct a gap assessment to compare the organization's current security stance to the detailed recommendations of the critical controls
- 2) Implement the “quick win” critical controls to address the gaps identified by the assessment over the next one or two quarters
- 3) Assign security personnel to analyze and understand how critical controls beyond the quick wins can be deployed in the organization's environment
- 4) Devise detailed plans to implement the "visibility and attribution" and "hardened configuration and improved information security hygiene" critical controls over the next year
- 5) Plan for deployment of the “advanced” controls” over the longer term.

Appendix A: Mapping between the 20 Critical Security Controls and National Institute of Standards and Technology Special Publication 800-53, Revision 3, Priority 1 Items

This mapping relates the controls set forth in this document to NIST Special Publication 800-53 Revision 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: Continuous Vulnerability Assessment and Remediation	RA-3 (a, b, c, d), RA-5 (a, b, 1, 2, 5, 6)
Critical Control 5: Malware Defenses	SC-18, SC-26, SI-3 (a, b, 1, 2, 5, 6)
Critical Control 6: Application Software Security	CM-7, RA-5 (a, 1), SA-3, SA-4 (3), SA-8, SI-3, SI-10
Critical Control 7: Wireless Device Control	AC-17, AC-18 (1, 2, 3, 4), SC-9 (1), SC-24, SI-4 (14, 15)
Critical Control 8: Data Recovery Capability	CP-9 (a, b, d, 1, 3), CP-10 (6)
Critical Control 9: Security Skills Assessment and Appropriate Training to Fill Gaps	AT-1, AT-2 (1), AT-3 (1)
Critical Control 10: 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 11: 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 12: Controlled Use of Administrative Privileges	AC-6 (2, 5), AC-17 (3), AC-19, AU-2 (4)
Critical Control 13: 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 14: 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 15: Controlled Access Based on the Need to Know	AC-1, AC-2 (b, c), AC-3 (4), AC-4, AC-6, MP-3, RA-2 (a)
Critical Control 16: Account Monitoring and Control	AC-2 (e, f, g, h, j, 2, 3, 4, 5), AC-3
Critical Control 17: 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 18: Incident Response Capability	IR-1, IR-2 (1), IR-4, IR-5, IR-6 (a), IR-8
Critical Control 19: Secure Network Engineering	IR-4 (2), SA-8, SC-7 (1, 13), SC-20, SC-21, SC-22, PM-7,
Critical Control 20: Penetration Tests and Red Team Exercises	CA-2 (1, 2), CA-7 (1, 2), RA-3, RA-5 (4, 9), SA-12 (7)

Appendix B: Mapping between the 20 Critical Security Controls and Australian Government's DSD 35 Mitigation Strategies

The following table maps the Australian Government's Defence Signals Directorate (DSD) top 35 strategies to mitigate targeted cyber intrusions to the Top 20 Critical Controls.

Mitigation Strategy Effectiveness Ranking	Mitigation Strategy	Matching Top 20 Critical Controls
1	Patch applications e.g. PDF viewer, Flash Player, Microsoft Office and Java. Patch or mitigate within two days for high risk vulnerabilities. Use the latest version of applications.	4.3
2	Patch operating system vulnerabilities. Patch or mitigate within two days for high risk vulnerabilities. Use the latest operating system version.	4.3
3	Minimise the number of users with domain or local administrative privileges. Such users should use a separate unprivileged account for email and web browsing.	19.1, 19.6
4	Application whitelisting to help prevent malicious software and other unapproved programs from running e.g. by using Microsoft Software Restriction Policies or AppLocker.	2.4
5	Host-based Intrusion Detection/Prevention System to identify anomalous behaviour such as process injection, keystroke logging, driver loading and call hooking.	8.1, 8.6
6	Whitelisted email content filtering allowing only attachment types required for business functionality. Preferably convert/sanitise PDF and Microsoft Office attachments.	8.5
7	Block spoofed emails using Sender Policy Framework checking of incoming emails, and a "hard fail" SPF record to help prevent spoofing of your organisation's domain.	12.5
8	User education e.g. Internet threats and spear phishing socially engineered emails. Avoid: weak passphrases, passphrase reuse, exposing email addresses, unapproved USB devices.	19.1, 17.1, 17.2, 17.3, 17.4, 17.5
9	Web content filtering of incoming and outgoing traffic, using signatures, reputation ratings and other heuristics, and whitelisting allowed types of web content.	12.1, 12.2, 12.3
10	Web domain whitelisting for all domains, since this	12.1, 12.7

	approach is more proactive and thorough than blacklisting a tiny percentage of malicious domains.	
11	Web domain whitelisting for HTTPS/SSL domains, since this approach is more proactive and thorough than blacklisting a tiny percentage of malicious domains.	12.1, 12.7
12	Workstation inspection of Microsoft Office files for abnormalities e.g. using the Microsoft Office File Validation feature.	8.1, 8.6
13	Application based workstation firewall, configured to deny traffic by default, to protect against malicious or otherwise unauthorised incoming network traffic.	3.3, 8.1, 5.1
14	Application based workstation firewall, configured to deny traffic by default, that whitelists which applications are allowed to generate outgoing network traffic.	3.3, 8.1, 8.8, 5.1
15	Network segmentation and segregation into security zones to protect sensitive information and critical services such as user authentication and user directory information.	10.8, 12.6, 20.4, 11.1, 11.5
16	Multi-factor authentication especially implemented for when the user is about to perform a privileged action, or access a database or other sensitive information repository.	10.6, 19.11
17	Randomised local administrator passphrases that are unique and complex for all computers. Use domain group privileges instead of local administrator accounts.	19.1, 19.7
18	Enforce a strong passphrase policy covering complexity, length, and avoiding both passphrase reuse and the use of dictionary words.	19.1, 19.8, 13.7
19	Border gateway using an IPv6-capable firewall to prevent computers directly accessing the Internet except via a split DNS server, an email server, or an authenticated web proxy.	10.5, 12.7, 11.3
20	Data Execution Prevention using hardware and software mechanisms for all software applications that support DEP.	3.3
21	Antivirus software with up to date signatures, reputation ratings and other heuristic detection capabilities. Use gateway and desktop antivirus software from different vendors.	8.1, 8.2, 8.5, 8.6
22	Non-persistent virtualised trusted operating environment with limited access to network file shares, for risky activities such as reading email and web browsing.	2.6
23	Centralised and time-synchronised logging of allowed and blocked network activity, with regular log analysis, storing logs for at least 18 months.	7.1, 7.3, 7.5, 7.6, 7.7

24	Centralised and time-synchronised logging of successful and failed computer events, with regular log analysis, storing logs for at least 18 months.	7.1, 7.4, 7.5, 7.6
25	Standard Operating Environment with unrequired operating system functionality disabled e.g. IPv6, autorun and Remote Desktop. Harden file and registry permissions.	3.1, 3.2, 3.3, 8.3
26	Workstation application security configuration hardening e.g. disable unrequired features in PDF viewers, Microsoft Office applications, and web browsers.	3.1, 3.2, 3.3
27	Restrict access to NetBIOS services running on workstations and on servers where possible.	20.3, 20.4
28	Server application security configuration hardening e.g. databases, web applications, customer relationship management and other data storage systems.	3.1, 3.2, 3.3
29	Removable and portable media control as part of a Data Loss Prevention strategy, including storage, handling, whitelisting allowed USB devices, encryption and destruction.	8.3, 8.4, 9.7, 9.8, 9.10
30	TLS encryption between email servers to help prevent legitimate emails being intercepted and used for social engineering. Perform content scanning after email traffic is decrypted.	20.4
31	Disable LanMan password support and cached credentials on workstations and servers, to make it harder for adversaries to crack password hashes.	3.1, 3.2, 3.3, 19.5
32	Block attempts to access web sites by their IP address instead of by their domain name.	12.1, 12.7
33	Network-based Intrusion Detection/Prevention System using signatures and heuristics to identify anomalous traffic both internally and crossing network perimeter boundaries.	12.2, 12.3
34	Gateway blacklisting to block access to known malicious domains and IP addresses, including dynamic and other domains provided free to anonymous Internet users.	12.1
35	Full network traffic capture to perform post-incident analysis of successful intrusions, storing network traffic for at least the previous seven days.	12.4

Appendix C: Attack Types

As described in the introduction, 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 distribute hostile content on Internet-accessible (and sometimes internal) websites that exploits unpatched and improperly secured client software running on victim machines.	2, 3
Attackers continually scan for vulnerable software and exploit it to gain control of target machines.	2, 4
Attackers use currently infected or compromised machines to identify and exploit other vulnerable machines across an internal network.	2, 10
Attackers exploit weak default configurations of systems that are more geared to ease of use than security.	3
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.	4
Attackers compromise target organizations that do not exercise their defenses to determine and continually improve their effectiveness.	4, 20
Attackers use malicious code to gain and maintain control of target machines, capture sensitive data, and then spread it to other systems, sometimes wielding code that disables or dodges signature-based anti-virus tools.	5

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.	5, 10
Attackers exploit weak application software, particularly web applications, through attack vectors such as SQL injection, cross-site scripting, and similar tools.	6
Attackers exploit wireless access points to gain entry into a target organization's internal network, and exploit wireless client systems to steal sensitive information.	7
Attackers exploit users and system administrators via social engineering scams that work because of a lack of security skills and awareness.	9
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.	10, 13
Attackers trick a user with an administrator level account into opening a phishing-style e-mail with an attachment or surfing to the attacker's content on an Internet website, allowing the attacker's malicious code or exploit to run on the victim machine with full administrator privileges.	12
Attackers exploit boundary systems on Internet-accessible DMZ networks, and then pivot to gain deeper access on internal networks.	13
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.	13, 19
Attackers operate undetected for extended periods of time on compromised systems because of a lack of logging and log review.	14
Attackers gain access to sensitive documents in an organization that does not properly identify and protect sensitive information or separate it from non-sensitive information.	15
Attackers compromise inactive user accounts left behind by temporary workers, contractors, and former employees, including accounts left behind by the	16

attackers themselves who are former employees.	
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.	16
Attackers gain access to internal enterprise systems gather and exfiltrate sensitive information without detection by the victim organization.	17
Attackers compromise systems and alter important data, potentially jeopardizing organizational effectiveness via polluted information.	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 attacker's presence, or recover to a secure production state.	18