Critical Controls for Effective Cyber Defense

Table of Contents

Introduction

Critical Controls for Effective Cyber Defense

Version 4.0

To secure organizations from cyber attacks, networks and systems must vigorously defend against a variety of internal and external threats. Defenses must also be prepared to detect and thwart follow-on attacks inside a network that has already been compromised. Two guiding principles are: "Prevention is ideal but detection is a must" and "Offense informs defense."

The Goal

The goal of the Critical Controls is to strengthen the defensive posture of your organization's information security; reduce compromises, recovery efforts, and associated costs; and protect critical assets and infrastructure. The Controls provide continuous, automated monitoring of the riskiest portions of your information technology infrastructure. Having them in place will allow your organization to focus on its primary mission.

The Methodology

The Critical Controls provide a prioritized, risk-based approach to security based on actual threats. The Controls focus on automation to provide cost efficiency, measurable results, scalability, and reliability. The four critical tenets of an effective cyber defense system as reflected in the Critical Controls are:

- Offense Informs Defense: Use knowledge of actual attacks that have compromised systems to provide the foundation to build effective defenses.
- Metrics: Establish common metrics to provide a shared language for executives, IT specialists, auditors, and security officials to measure the effectiveness of security measures within an organization so that required adjustments can be identified and implemented quickly.
- Continuous monitoring: Carry out continuous monitoring/auditing to test and validate whether current security measures are proactively remediating vulnerabilities in a timely manner.
- Automation: Automate defenses so that organizations can achieve reliable, scalable, and continuous measurements of their adherence to the controls and related metrics.

Why the Controls Work

The Critical Controls draw on the knowledge gained in combating the myriad attacks launched regularly against networks. Top cybersecurity experts have joined forces to make the Controls the most effective and specific set of technical measures available to detect and prevent the most common and damaging of those computer attacks. This consensus process is the foundation of the Controls because it provides first-hand knowledge of actual attacks and the best defensive

techniques to stop them. It also ensures that the Controls will address the root causes of attacks so that security measures deployed today will be effective against the next generation of advanced threats.

The Critical Controls represent the sum total of efforts over the last decade to develop standards to identify common vulnerabilities and their severity, define secure configurations, inventory systems and platforms, and pinpoint application weaknesses. These efforts include the Associated Manageable Network Plan Milestones and Network Security Tasks developed by the National Security Agency (NSA), and the Security Content Automation Program sponsored by the National Institute of Standards and Technology (NIST).

How to Apply the Controls

The Critical Controls are 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 controls in place. To help organizations with different levels of information security capabilities design a sound security baseline and then improve beyond that, the sub-controls included in each of the Critical Control summaries specify actions that organizations can take to improve their defenses:

- *Quick wins* on fundamental aspects of information security to help an organization rapidly improve its security stance without major procedural, architectural, or technical changes to its environment.
- *Visibility and attribution measures* to improve the process, architecture, and technical capabilities of organizations to monitor their networks and computer systems to detect attack attempts, locate points of entry, identify already-compromised machines, interrupt infiltrated attackers' activities, and gain information about the sources of an attack.
- *Improved information security configuration and hygiene* to reduce the number and magnitude of security vulnerabilities and improve the operations of networked computer systems, with a focus on protecting against poor security practices by system administrators and end-users that could give an attacker an advantage.
- *Advanced sub-controls* that use new technologies that provide maximum security but are harder to deploy or more expensive than commoditized security solutions.

Prioritizing the Controls

As a result of the rigorous process of analysis, debate, and consensus that is the basis of their design, the Critical Controls are categorized by their attack migration impact and their importance in providing a direct defense against attacks. The rankings can be tailored to the cybersecurity needs and capabilities of a particular organization. Priority is given to Controls that (1) mitigate known attacks, (2) address a wide variety of attacks, and (3) identify and stop attacks early in the compromise cycle.

Contributors

The strength of the Critical Controls is in the many organizations that contributed input to identifying and designing them. These include the U.S. Departments of Defense and Energy, the U.S. Computer Emergency Readiness Team, the FBI and other law enforcement agencies, and civilian penetration testers.

The Critical Controls Document

The presentation of each Critical Control in this document includes:

- A step-by-step breakdown of the procedures and tools required to implement and automate it;
- An explanation of how attackers exploit the absence of this control;
- Entity relationship diagrams to show how the controls can be implemented;
- An outline of the most appropriate sub-controls to implement, automate, and measure effectiveness;
- Summaries of metrics and tests that can be used to evaluate implementation; and
- A list of associated NIST controls and NSA tasks.

After organizations implement the Controls and gain experience with automation, CIOs can use the document as an audit guide to ensure that they are taking the right actions for effective cyber defense, and IGs can use it to verify the CIOs' tests.

The Controls are meant to deal with multiple kinds of computer attackers, including 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. The Controls are not limited to blocking 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 reducing the initial attack surface by hardening security, identifying compromised machines to address long-term threats inside an organization's network, controlling super-user privileges on systems, and disrupting attackers' command-and-control of implanted malicious code.

Finally, each control included in this document describes a series of tests that organizations can conduct on a periodic or continual basis to ensure that appropriate defenses are in place.

Moving Ahead

The consensus effort to define critical security controls is an evolving process. Changing technology and attack patterns will necessitate future changes to the current set of Critical Controls. In a sense, this will be a living document moving forward, but the Controls described here are a solid start toward making fundamental computer security defenses a well-understood, replicable, measurable, scalable, and reliable process.

Description of Controls

Critical Control 1: Inventory of Authorized and Unauthorized Devices

The processes and tools used to track/control/prevent/correct network access by devices (computers, network components, printers, anything with IP addresses) based on an asset inventory of which devices are allowed to connect to the network.

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.

APTs (advanced persistent threat) target internal users with the goal of compromising a system on the private network that can be used as a pivot point to attack internal systems. Even systems that are connected to the private network, without visibility from the Internet, can still be a target of the advanced adversary. Any system, even test systems that are connected for a short period of time, can still be used as a relay point to cause damage to an organization.

As new technology continues to come out, BYOD (bring your own device)— where employees bring personal devices into work and connect them to the network—is becoming very common. These devices could already be compromised and be used to infect internal resources.

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

1. *Quick wins*: Deploy an automated asset inventory discovery tool and use it to build a preliminary asset inventory of systems connected to an organization's public and private network(s). Both active tools that scan through network address ranges and passive tools that identify hosts based on analyzing their traffic should be employed.
2. *Quick wins:* Deploy DHCP Server logging, and utilize a system to improve the asset inventory and help detect unknown systems through this DHCP information.
3. *Quick wins:* All equipment acquisitions should automatically update the inventory system as new, approved devices are connected to the network. A robust change control process can also be used to validate and approve all new devices.
4. *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, Voice Over-IP telephones, multi-homed addresses, virtual addresses, etc. The asset inventory created must also include data on whether the device is a portable and/or personal 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.

5. *Configuration/Hygiene:* Make sure the asset inventory database is properly protected and a copy stored in a secure location.
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. *Configuration/Hygiene:* Deploy network access control (NAC) 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.
9. *Configuration/Hygiene:* Create separate VLANs for BYOD (bring your own device) systems or other untrusted devices.
10. *Advanced:* Utilize client certificates to validate and authenticate systems prior to connecting to the private network.

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 the Network
Milestone 3: Network Architecture
Network Access Protection/Control (NAP/NAC)

Procedures and Tools to Implement and Automate this Control

Organizations must first establish information/asset owners, deciding and documenting which organizations and individuals are responsible for each component of a business process that includes information, software, and hardware. 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 it. In particular, when organizations acquire new systems, they record the owner and features of each new asset, including its network interface media access control (MAC) address and location. 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 and NAC to only allow authorized systems that are properly configured to connect to the network.

Going further, effective organizations configure free or commercial network scanning tools to perform network sweeps on a regular basis, 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. 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 using an IP address 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 the unauthorized system 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 automated tools, notification about an unauthorized asset connected to the network can be sent within two minutes and isolation achieved 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 System Entity Relationship Diagram (ERD)

Organizations will find that by diagramming the entities necessary to fully meet the goals defined in this control, it will be easier to identify how to implement them, test the controls, and identify where potential failures in the system might occur.

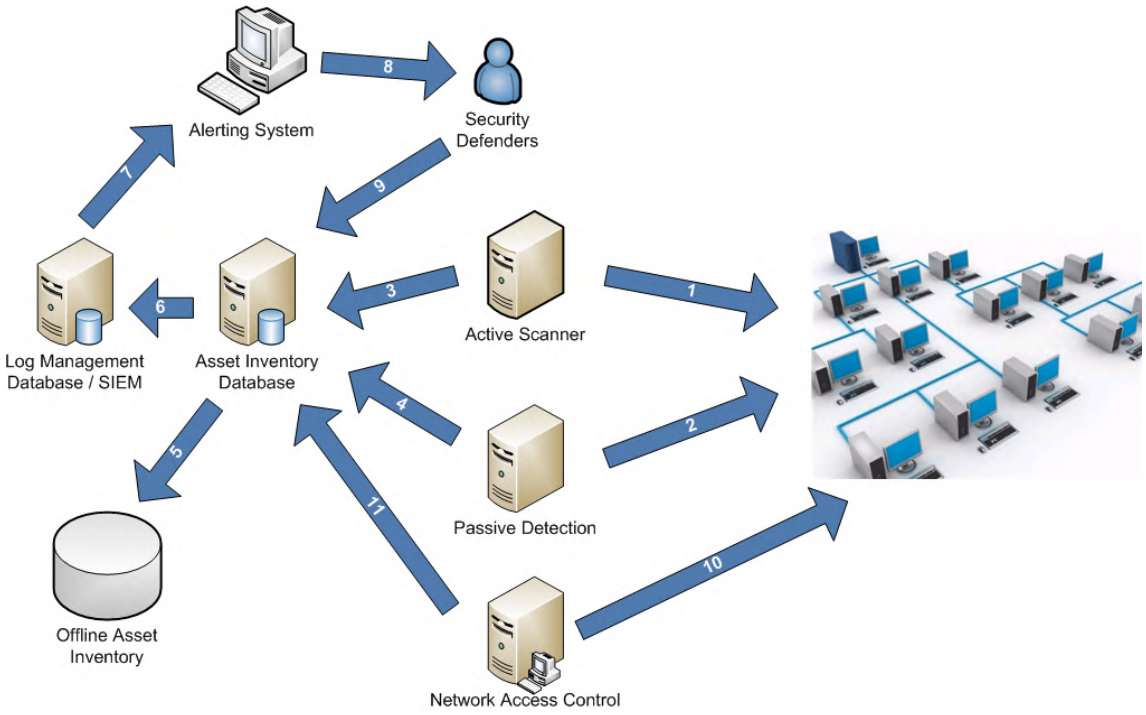

A control system is a device or set of devices used to manage, command, direct, or regulate the behavior of other devices or systems. In this case, we are examining hardware devices on the organization's network. These systems should be able to identify if new systems are introduced to the environment that have not been authorized by enterprise personnel. The following list of the steps in the above diagram shows how the entities work together to meet the business goal defined in this control. The list also helps identify each of the process steps in order to help identify potential failure points in the overall control.

Step 1: Active device scanner scans network systems

Step 2: Passive device scanner captures system information

Step 3: Active scanner reports to inventory database

Step 4: Passive scanner reports to inventory database

Step 5: Inventory database stored offline

Step 6: Inventory database initiates alerting system

Step 7: Alert system notifies security defenders

Step 8: Security defenders monitor and secure inventory database

Step 9: Security defenders updates secure inventory database

Step 10: Network access control (NAC) continuously monitors network

Step 11: Network access control checks and provides updates to the asset inventory database

Critical Control 2: Inventory of Authorized and Unauthorized Software

The processes and tools organizations use to track/control/prevent/correct installation and execution of software on computers based on an asset inventory of approved 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. This list should be tied to file integrity checking software to validate that the software has not be modified.
2. *Quick wins:* Perform regular scanning and generate alerts when unapproved software is installed on a computer. A strict change control process should also be implemented to control any changes or installation of software to any systems on the network.
3. *Visibility/Attribution:* Deploy application white listing technology that allows systems to run only approved software and prevents execution of all other software on the system.
4. *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. The software inventory should be tied to vulnerability reporting/threat intelligence services to fix vulnerable software proactively.

5. *Visibility/Attribution:* The software inventory systems must be tied into the hardware asset inventory so all devices and associated software are tracked from a single location.
6. *Configuration/Hygiene:* 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. Dangerous file types (e.g., exe, zip, msi, etc.) should be closely monitored and/or blocked.
7. *Configuration/Hygiene:* Software inventory and application white listing should also be deployed on all mobile devices that are utilized across the organization.
8. *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.
9. *Advanced:* Configure client workstations with nonpersistent virtualized operating environments that can be quickly and easily restored to a trusted snapshot on a periodic basis.
10. *Advanced:* Deploy software that only provides signed software ID tags. A software identification tag is an XML file that is installed alongside software, and uniquely identifies the software, providing data for software inventory and asset management.

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 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 the unauthorized system 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.

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 System Entity Relationship Diagram (ERD)

Organizations will find that by diagramming the entities necessary to fully meet the goals defined in this control, it will be easier to identify how to implement them, test the controls, and identify where potential failures in the system might occur.

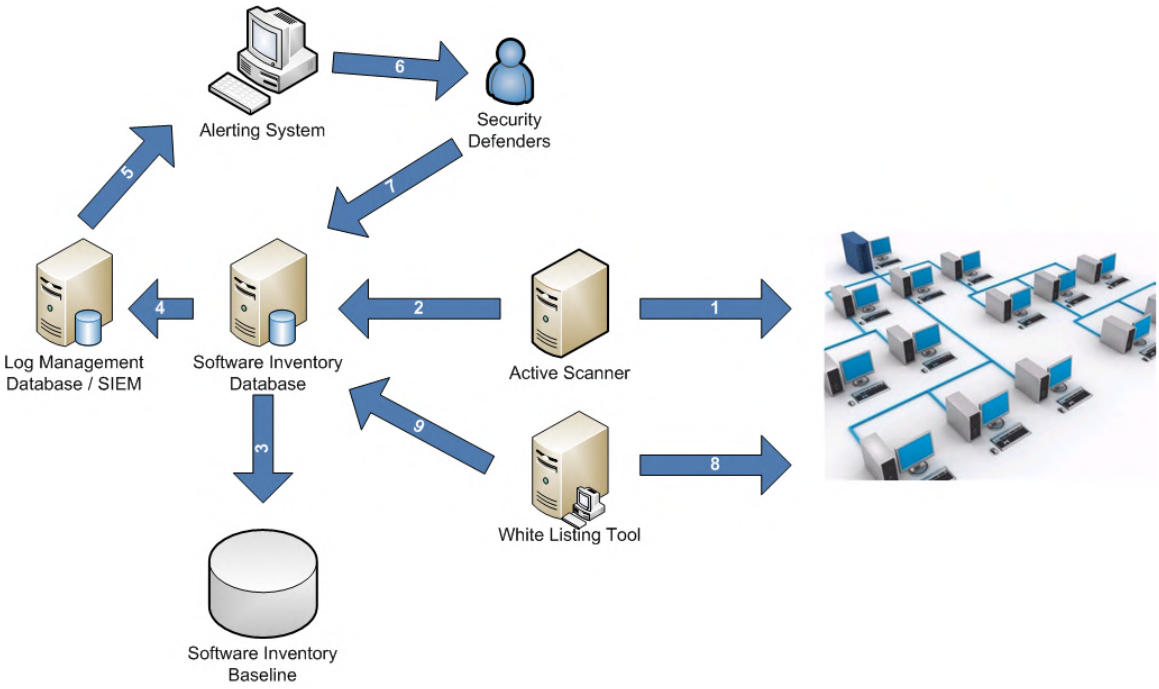

A control system is a device or set of devices to manage, command, direct, or regulate the behavior of other devices or systems. In this case, we are examining software installed on the organization's network systems. These systems should be able to identify if new software is introduced to the environment that has not been authorized by enterprise personnel. The following list of the steps in the above diagram shows how the entities work together to meet the business goal defined in this control. The list also helps identify what each of the process steps is in order to help identify potential failure points in the overall control.

Step 1: Active device scanner

Step 2: Active scanner reports to inventory database

Step 3: Inventory database compares to inventory baseline

Step 4: Inventory database initiates alerting system

Step 5: Alert system notifies security defenders

Step 6: Security defenders monitor and secure inventory database.

Step 7: Security defenders updates software inventory database

Step 8: White listing tool continuously monitors all systems on the network

Step 9: White listing checks and makes updates to the software inventory database

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

The processes and tools organizations use to track/control/prevent/correct security weaknesses in the configurations of the hardware and software of mobile devices, laptops, workstations, and servers based on a formal configuration management and change control process.

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. In addition, patches are not always applied in a timely manner and software updates often introduce unknown weaknesses into a piece of software that can be exploited by zero-day exploits. 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. Images should be created for both workstations and servers.
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 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 nonexecutable stacks and heaps. 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:* 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. Images should be tested at the hot or warm disaster recovery site if one is available.

5. *Quick wins:* Run the last version of software and make sure it is fully patched. Remove outdated or older software from the system.
6. *Visibility/Attribution:* Any deviations from the standard build or updates to the standard build should be approved by a change control board and documented in a change management system.
7. *Visibility/Attribution:* 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.
8. *Visibility/Attribution:* Utilize application white listing to control and manage any configuration changes to the software running on the system.
9. *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 others that do not actively support strong encryption should only be used if they are performed over a secondary encryption channel, such as SSL or IPSEC.
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 tools compliant with Security Content Automation Protocol (SCAP) 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:* Deploy system configuration management tools, such as Active Directory Group Policy Objects for Microsoft Windows systems or Puppet for Unix systems that will automatically enforce and redeploy configuration settings to systems at regularly scheduled intervals.
13. *Advanced:* Organizations need to adopt a formal process and management infrastructure for configuration control of mobile devices. The process needs to include secure remote wiping of lost or stolen devices, approval of corporate apps, and denial of unapproved apps. If the device is owned by the organization, a full wipe should be performed. If it is a BYOD system, a selective wipe should be performed, removing the organization's information.

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 a combination 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 the unauthorized system 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.

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 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 notification process.
2. System scanning tools that check for 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 System Entity Relationship Diagram (ERD)

Organizations will find that by diagramming the entities necessary to fully meet the goals defined in this control, it will be easier to identify how to implement the, test the controls, and identify where potential failures in the system might occur. As with any configurations, all changes must be approved and managed by a change control process.

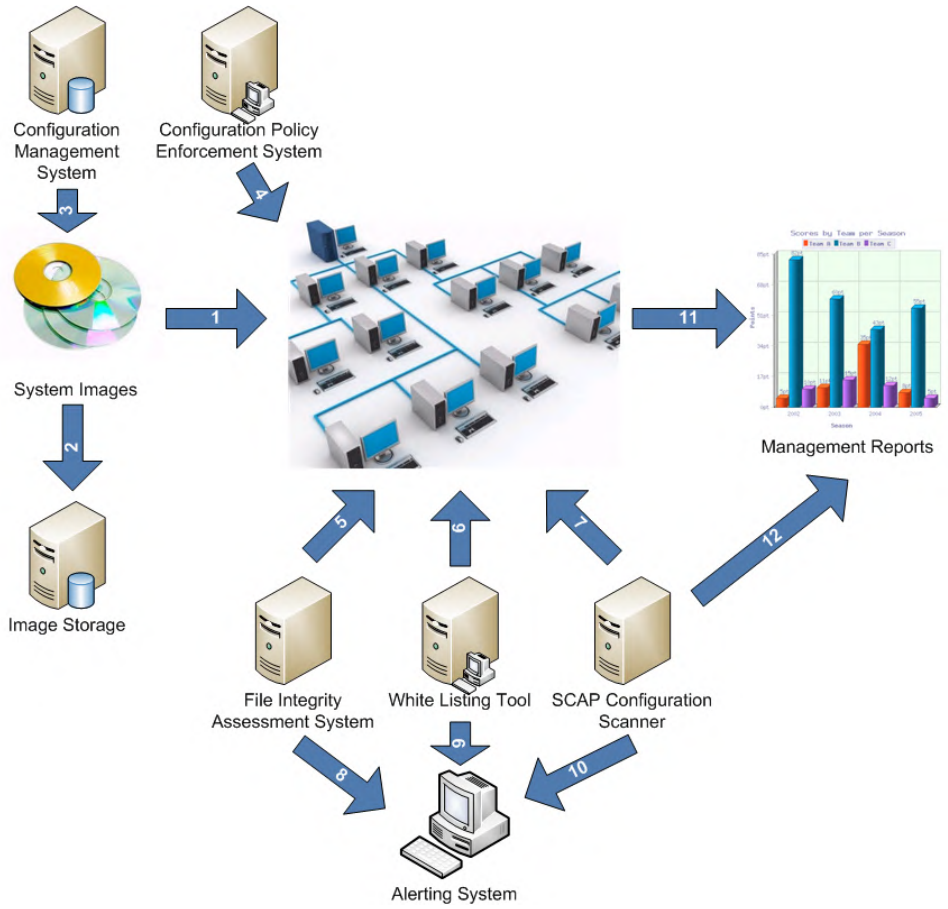

A control system is a device or set of devices to manage, command, direct, or regulate the behavior of other devices or systems. In this case, we are examining the devices, software, and entities used to manage and implement consistent configuration settings to workstations, laptops, and servers on the network. We will be discussing other network devices later in the course. The following list of the steps in the above diagram shows how the entities work together to meet the business goal defined in this control. The list also helps identify each step in order to identify potential failure points in the overall control.

Step 1: Secured system images applied to computer systems

Step 2: Secured system images stored in a secure manner

Step 3: Configuration management system validates and checks system images

Step 4: Configuration policy enforcement system actively scans production systems for misconfigurations or deviations from baselines

Step 5: File integrity assessment systems monitor critical system binaries and data sets

Step 6: White listing tool monitors systems configurations and software

Step 7: SCAP configuration scanner validates configurations

Step 8: File integrity assessment system sends deviations to alerting system

Step 9: White listing tool sends deviations to alerting system

Step 10: SCAP configuration scanner sends deviations to alerting system

Step 11 and 12: Management reports document configuration status

Critical Control 4: Continuous Vulnerability Assessment and Remediation

The processes and tools used to detect/prevent/correct security vulnerabilities in the configurations of devices that are listed and approved in the asset inventory database.

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. Vulnerabilities must also be tied to threat intelligence and be properly prioritized.

As vulnerability scans become more common, attackers are utilizing them as a point of exploitation. It is important to carefully control authenticated vulnerability scans and the associated administrator account. Attackers will take over one machine with local privileges, and wait for an authenticated scan to occur against the machine. When the scanner logs in with domain admin privileges, the attacker either grabs the token of the logged-in scanning tool, or sniffs the challenge response and cracks it. Either way, the attacker then can pivot anywhere else in the organization as domain admin.

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

1. *Quick wins:* Run automated vulnerability scanning tools against all systems on their networks on a weekly or more frequent basis using a SCAP-validated vulnerability scanner that looks for both code-based vulnerabilities (CVE) and configuration-based vulnerabilities (CCE). 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 target known to be vulnerable.
3. *Quick wins:* Utilize a dedicated account for authenticated vulnerability scans. The scanning account should not be used for any other administrative activities and tied to specific IP addresses. Ensure only authorized employees have access to the vulnerability management user interface and that roles are applied to each user.
4. *Quick wins:* Subscribe to vulnerability intelligence services in order to stay aware of emerging exposures.
5. *Visibility/Attribution:* Deploy automated patch management tools and software update tools for operating system and software/applications on all systems for which such tools

are available and safe. Patches should be applied to all systems, even systems that are properly air gapped.

6. *Visibility/Attribution:* Carefully monitor logs associated with any scanning activity and associated administrator accounts to ensure that all scanning activity and associated access via the privileged account is limited to the timeframes of legitimate scans.
7. *Configuration/Hygiene:* In addition to 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.
8. *Configuration/Hygiene:* 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.
9. *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.
10. *Configuration/Hygiene:* 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. Alternative countermeasures should be considered if patches are not available.
11. *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.
12. *Configuration/Hygiene:* Address the most damaging vulnerabilities first. Prioritize the vulnerable assets based on both the technical and organization-specific business risks. An industry-wide or corporate-wide vulnerability ranking may be inadequate to prioritize which specific assets to address first. A phased rollout can be used to minimize the impact to the organization.

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. 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. 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 System Entity Relationship Diagram (ERD)

Organizations will find that by diagramming the entities necessary to fully meet the goals defined in this control, it will be easier to identify how to implement tem, test the controls, and identify where potential failures in the system might occur.

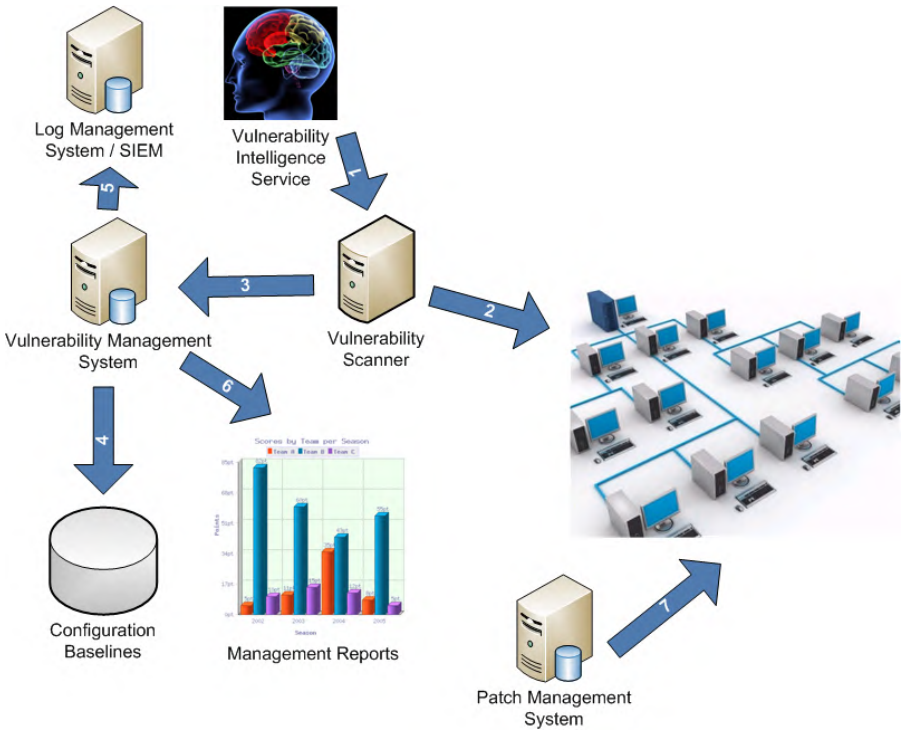

A control system is a device or set of devices to manage, command, direct, or regulate the behavior of other devices or systems. In this case, the vulnerability scanners, management system, patch management systems, and configuration baselines all work together to address an organization's vulnerability management and remediation strategy. The following list of the steps in the above diagram shows how the entities work together to meet the business goal defined in this control. The list also helps identify what each of the process steps is in order to help identify potential failure points in the overall control.

Step 1: Vulnerability intelligence service provides inputs to vulnerability scanner

Step 2: Vulnerability scanners scan production systems

Step 3: Vulnerability scanners report detected vulnerabilities to a vulnerability management system (VMS)

Step 4: The VMS compares production systems to configuration baselines

Step 5: The VMS sends information to log management correlation system

Step 6: The VMS produces reports for management

Step 7: A patch management system applies software updates to production systems

Critical Control 5: Malware Defenses

The processes and tools used to detect/prevent/correct installation and execution of malicious software on all devices.

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, the cloud, 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:* 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. The endpoint security solution should include zero-day protection such as network behavioral heuristics.
2. *Quick wins:* 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:* 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. If the devices are not required for business use, they should be disabled.
4. *Quick wins:* Configure systems so that they conduct an automated anti-malware scan of removable media when it is inserted.
5. *Quick wins:* All e-mail 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 e-mail content filtering and web content filtering.
6. *Quick wins:* Apply anti-virus scanning at the Web Proxy gateway. Content filtering for file-types should be applied at the perimeter.
7. *Quick wins:* Deploy features and toolkits such as Data Execution Prevention (DEP) and Enhanced Mitigation Experience Toolkit (EMET), products that provide sandboxing (e.g., run browsers in a VM), and other techniques that prevent malware exploitation.
8. *Quick wins:* Limit use of external devices to those that have business need. Monitor for use and attempted use of external devices.

9. *Visibility/Attribution:* Block access to external e-mail systems, instant messaging services and other social media tools.
10. *Visibility/Attribution:* Automated monitoring tools should use behavior-based anomaly detection to complement and enhance traditional signature-based detection.
11. *Visibility/Attribution:* Utilize network-based anti-malware tools to analyze all inbound traffic and filter out malicious content before it arrives at the endpoint.
12. *Advanced:* Continuous monitoring should be performed on all inbound and outbound traffic. Any large transfers of data or unauthorized traffic should be flagged and, if validated as malicious, the computer should be moved to an isolated VLAN.
13. *Advanced:* Implement an incident response process that 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 security vendor for "out-of-band" signature creation and deployed to the enterprise by system administrators.
14. *Advanced:* Utilize network-based flow analysis tools to analyze inbound and outbound traffic looking for anomalies, indicators of malware, and compromised systems.
15. *Advanced:* Deploy "reputation-based technologies" on all endpoint devices to cover the gap of signature based technologies.
16. Advanced: Enable domain name system (DNS) query logging to detect hostname lookup for known malicious C2 domains.
17. *Advanced:* Apply proxy technology to all communication between internal network and the Internet.

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, organizations use automation. They use the built-in administrative features of enterprise endpoint 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.

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.

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 System Entity Relationship Diagram (ERD)

Organizations will find that by diagramming the entities necessary to fully meet the goals defined in this control, it will be easier to identify how to implement them, test the controls, and identify where potential failures in the system might occur.

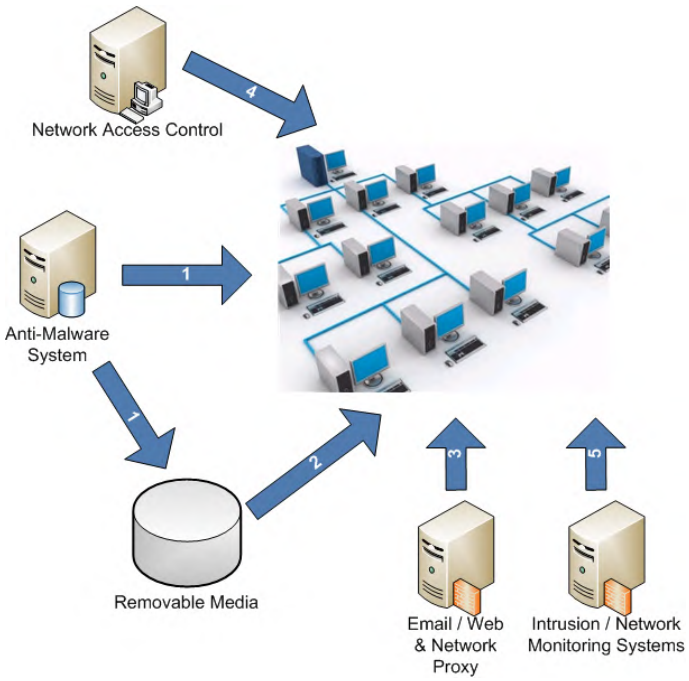

A control system is a device or set of devices to manage, command, direct, or regulate the behavior of other devices or systems. In this case, we are examining anti-malware systems and threat vectors such as removable media. The following list of the steps in the above diagram shows how the entities work together to meet the business goal defined in this control. The list also helps identify what each of the process steps is in order to help identify potential failure points in the overall control.

Step 1: Anti-malware systems analyzes production systems and removable media

Step 2: Removable media is analyzed when connected to production systems

Step 3: Email/web and network proxy devices analyze all incoming and outgoing traffic

Step 4: Network access control monitors all systems connected to the network

Step 5: Intrusion / network monitoring systems performs continuous monitoring looking for signs of malware

Critical Control 6: Application Software Security

The processes and tools organizations use to detect/prevent/correct security weaknesses in the development and acquisition of software applications.

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:* Protect web applications by deploying web application firewalls (WAFs) 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. *Visibility/Attribution:* 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. Organizations should understand how their applications behave under denial of service or resource exhaustion attacks.

4. *Visibility/Attribution:* System error messages should not be displayed to end-users (output sanitization).
5. *Visibility/Attribution:* Maintain separate environments for production and nonproduction systems. Developers should not typically have unmonitored access to production environments.
6. *Configuration/Hygiene:* 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.
7. *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. All systems that are part of critical business processes should also be tested.
8. *Configuration/Hygiene:* Ensure that all software development personnel receive training in writing secure code for their specific development environment.
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. 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, 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 been unsuccessful. 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 System Entity Relationship Diagram (ERD)

Organizations will find that by diagramming the entities necessary to fully meet the goals defined in this control, it will be easier to identify how to implement them, test the controls, and identify where potential failures in the system might occur.

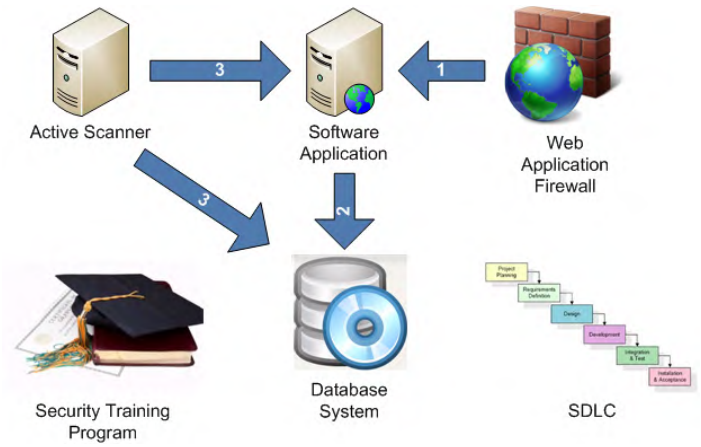

A control system is a device or set of devices used to manage, command, direct, or regulate the behavior of other devices or systems. In this case, we are examining the process of monitoring applications and using tools that enforce a security style when developing applications.

The following list of the steps in the above diagram shows how the entities work together to meet the business goal defined in this control. The list also helps identify what each of the process steps is in order to help identify potential failure points in the overall control.

Step 1: Web application firewalls protect connections to internal web applications

Step 2: Software applications securely connect to database systems

Step 3: Code analysis and vulnerability scanning tools scan application systems and database systems

Critical Control 7: Wireless Device Control

The processes and tools used to track/control/prevent/correct the security use of wireless LANS (local area networks), access points, and wireless client systems.

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 outside the physical building, bypassing organizations' security perimeters by connecting wirelessly to access points inside the organization. Wireless clients accompanying traveling 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:* 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:* 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:* 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 WIDS 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, 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, 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:* Ensure that all wireless traffic leverages at least Advanced Encryption Standard (AES) encryption used with at least WiFi Protected Access 2 (WPA2) protection.
10. *Configuration/Hygiene:* Ensure that wireless networks use authentication protocols such as Extensible Authentication Protocol-Transport Layer Security (EAP/TLS), which provide credential protection and mutual authentication.
11. *Configuration/Hygiene:* Ensure that wireless clients use strong, multi-factor authentication credentials to mitigate the risk of unauthorized access from compromised credentials.
12. *Configuration/Hygiene:* Disable peer-to-peer wireless network capabilities on wireless clients, unless such functionality meets a documented business need.
13. *Configuration/Hygiene:* Disable wireless peripheral access of devices (such as Bluetooth), unless such access is required for a documented business need.
14. *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.
15. *Configuration/Hygiene:* All mobile devices, including personnel devices, must be registered prior to connecting to the wireless network. All registered devices must be scanned and follow the corporate policy for host hardening and configuration management.
16. *Advanced:* Configure all wireless clients used to access private 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 organization.

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 System Entity Relationship Diagram (ERD)

Organizations will find that by diagramming the entities necessary to fully meet the goals defined in this control, it will be easier to identify how to implement them, test the controls, and identify where potential failures in the system might occur.

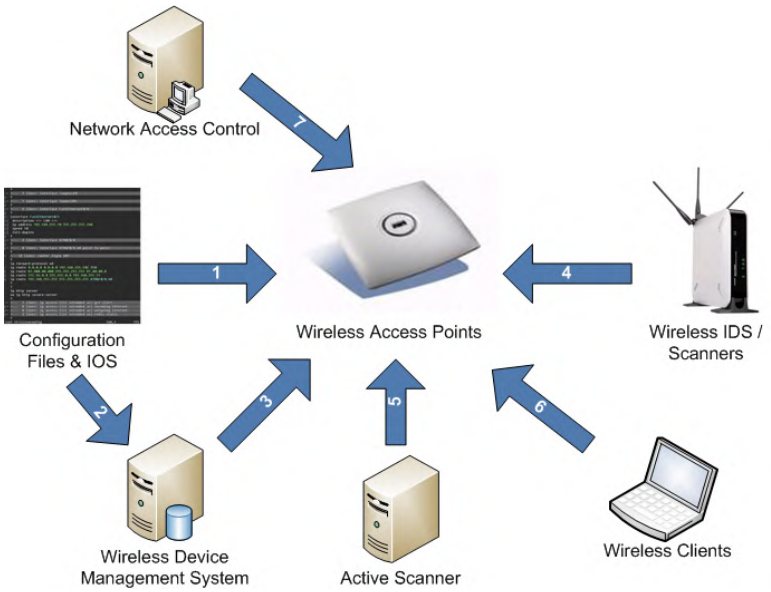

A control system is a device or set of devices used to manage, command, direct, or regulate the behavior of other devices or systems. In this case, we are examining the configuration and management of wireless devices, wireless IDS/scanners, wireless device management systems, and vulnerability scanners. The following list of the steps in the above diagram shows how the entities work together to meet the business goal defined in this control. The list also helps identify what each of the process steps is in order to help identify potential failure points in the overall control.

Step 1: Hardened configurations applied to wireless devices

Step 2: Hardened configurations managed by a configuration management system

Step 3: Configuration management system manages the configurations on wireless devices

Step 4: Wireless IDS monitor usage of wireless communications

Step 5: Vulnerability scanners scan wireless devices for potential vulnerabilities

Step 6: Wireless clients utilize wireless infrastructure systems in a secure manner

Critical Control 8: Data Recovery Capability

The processes and tools used to properly back up critical information with a proven methodology for timely recovery of it.

Note: This control has one or more sub-controls that must be validated manually.

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:* 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. All backup policies should be compliant with any regulatory or official requirements.
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 personal 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:* Ensure that backups are properly protected via physical security or encryption when they are stored, as well as when they are moved across the network. This includes remote backups and cloud services.
5. *Configuration/Hygiene:* Backup media, such as hard drives and tapes, should be stored in physically secure, locked facilities. End-of-life backup media should be securely erased/destroyed.

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 (or whenever new backup equipment is purchased), 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 System Entity Relationship Diagram (ERD)

Organizations will find that by diagramming the entities necessary to fully meet the goals defined in this control, it will be easier to identify how to implement them, test the controls, and identify where potential failures in the system might occur.

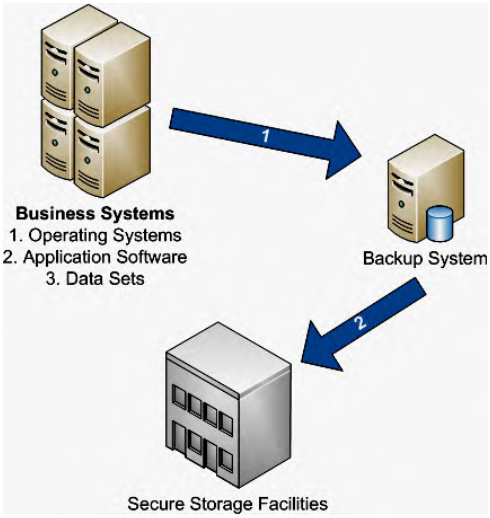

A control system is a device or set of devices used to manage, command, direct, or regulate the behavior of other devices or systems. In this case, we are examining an organization's capability to restore systems in the event that data needs to be restored because of a data loss or breach of a system. While backups are certainly an important part of this process, the ability to restore data is the critical component. The following list of the steps in the above diagram shows how the entities work together to meet the business goal defined in this control. The list also helps identify what each of the process steps is in order to help identify potential failure points in the overall control.

Step 1: Production business systems backed up on a regular basis to authorized organizational backup systems

Step 2: Backups created are stored offline at secure storage facilities.

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

The process and tools to make sure an organization understands the technical skill gaps within its workforce, including an integrated plan to fill the gaps through policy, training, and awareness.

Note: This control has one or more sub-controls that must be validated manually.

How Do Attackers Exploit the Absence of this Control?

The following are the 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 must 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:* Perform gap analysis to see which security areas employees are not adhering to and use this as the basis for an awareness program. Organizations should devise periodic security awareness assessments 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.

2. *Quick wins:* 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.
3. *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.
4. *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.
5. *Configuration/Hygiene:* 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. *Configuration/Hygiene:* 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 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 System Entity Relationship Diagram (ERD)

Organizations will find that by diagramming the entities necessary to fully meet the goals defined in this control, it will be easier to identify how to implement them, test the controls, and identify where potential failures in the system might occur.

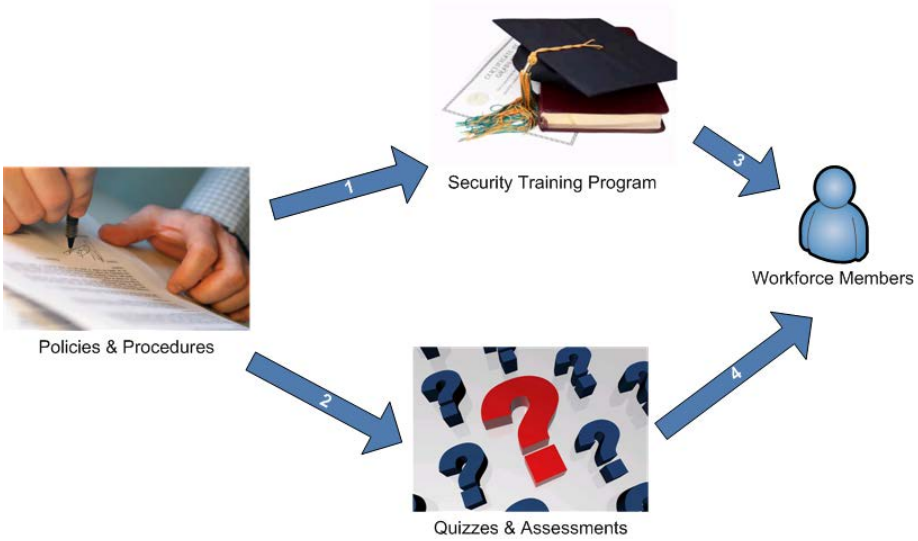

A control system is a device or set of devices used to manage, command, direct, or regulate the behavior of other devices or systems. In this case, we are examining the importance of educating workforce members in the security knowledge that they need in order to best perform their roles and responsibilities within the organization. The following list of the steps in the above diagram show how the entities work together to meet the business goal defined in this control. The list also helps identify what each of the process steps is in order to help identify potential failure points in the overall control.

Step 1: Create training and awareness programs about the organization's information security policies and procedures

Step 2: Design quizzes and assessments about the organization's information security policies and procedures

Step 3: Present the training and awareness program to the organization's workforce

Step 4: Present the quizzes and assessments to the organization's workforce and validate the understanding of the information presented in the training and awareness program

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

The processes and tools used to track/control/prevent/correct security weaknesses in the configurations in network devices such as firewalls, routers, and switches based on formal configuration management and change control processes.

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 and can change over time. 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. *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.
4. *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.

5. *Configuration/Hygiene:* Network devices should be managed using two-factor authentication and encrypted sessions.
6. *Configuration/Hygiene:* The latest stable version of any security-related updates must be installed within 30 days of the update being released from the device vendor.
7. *Advanced:* The network infrastructure should be managed across network connections that are separated from the business use of that network, relying on separate VLANs or, preferably, 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 alert 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 System Entity Relationship Diagram (ERD)

Organizations will find that by diagramming the entities necessary to fully meet the goals defined in this control, it will be easier to identify how to implement them, test the controls, and identify where potential failures in the system might occur.

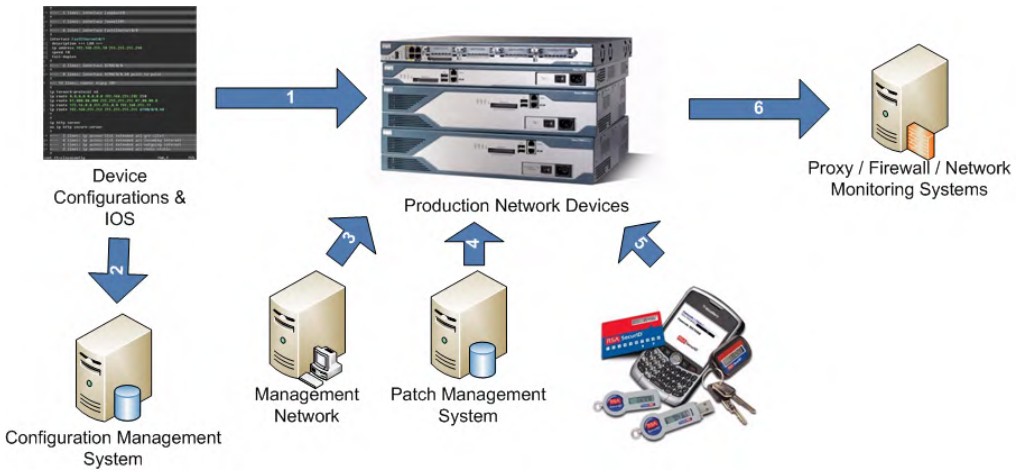

A control system is a device or set of devices used to manage, command, direct, or regulate the behavior of other devices or systems. In this case we are examining the network devices, test lab network devices, configuration systems, and configuration management devices. The following list of the steps in the diagram shows how the entities work together to meet the business goal defined in this control. The list also helps identify each step in order to help identify potential failure points in the overall control.

Step 1: Hardened device configurations applied to production devices

Step 2: Hardened device configuration stored in a secure configuration management system

Step 3: Management network system validates configurations on production network devices

Step 4: Patch management system applies tested software updates to production network devices

Step 5: Two-factor authentication system required for administrative access to production devices

Step 6: Proxy/firewall/network monitoring systems analyzes all connections to production network devices

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

The processes and tools used to track/control/prevent/correct use of ports, protocols, and services on networked devices.

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:* Any service that is not needed should be turned off for 30 days and after 30 days uninstalled from the system.
2. *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.
3. *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 change that is not listed on the organization's approved baseline is discovered, an alert should be generated and reviewed.
4. *Quick wins:* All services should be kept up to date and any unnecessary components uninstalled and removed from the system.
5. *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.
6. *Configuration/Hygiene:* Services needed for business use across the internal network should be reviewed quarterly via a change control group, and business units should rejustify the business use. Services that are turned on for projects or limited engagements should be turned off when they are no longer needed and properly documented.
7. *Configuration/Hygiene:* Operate critical services on separate physical or logical host machines, such as DNS, file, mail, web, and database servers.
8. *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. 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.

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 System Entity Relationship Diagram (ERD)

Organizations will find that by diagramming the entities necessary to fully meet the goals defined in this control, it will be easier to identify how to implement them, test the controls, and identify where potential failures in the system might occur.

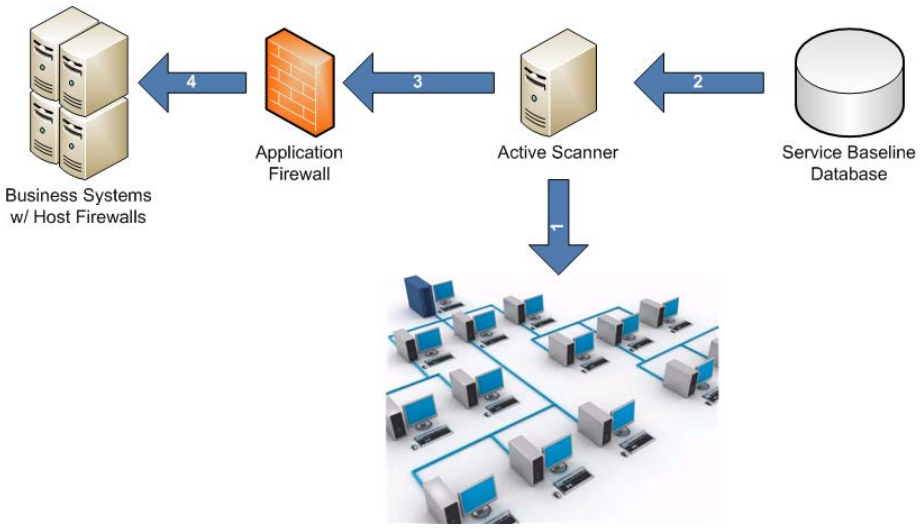

A control system is a device or set of devices used to manage, command, direct, or regulate the behavior of other devices or systems. In this case, we are examining how active scanning systems gather information on network devices and evaluate that data against the authorized service baseline database. The following list of the steps in the above diagram shows how the entities work together to meet the business goal defined in this control. The list also helps identify what each of the process steps is in order to help identify potential failure points in the overall control.

Step 1: Active scanner analyzes production systems for unauthorized ports, protocols, and services

Step 2: System baselines regularly updated based on necessary/required services

Step 3: Active scanner validates which ports, protocols and services are blocked or allowed by the application firewall

Step 4: Active scanner validates which ports, protocols and services are accessible on business systems protected with host based firewalls

Critical Control 12: Controlled Use of Administrative Privileges

The processes and tools used to track/control/prevent/correct the use, assignment, and configuration of administrative privileges on computers, networks, and applications.

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 common attack technique, 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.

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

1. *Quick wins:* 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.
2. *Quick wins:* Configure all administrative passwords to be complex and contain letters, numbers and special characters intermixed with no dictionary words present in the password. Strong passwords should be of a sufficient length to increase the difficultly it takes to crack the password. Pass phrases containing multiple dictionary words, along with special characters, are acceptable if they are of a reasonable length.
3. *Quick wins:* Configure all administrative-level accounts to require regular password changes on a frequent interval tied to the complexity of the password.
4. *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.
5. *Quick wins:* 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.
6. *Quick wins:* Passwords for all systems should be stored in a well-hashed or encrypted format, with weaker formats eliminated from the environment. Furthermore, files

containing these encrypted or hashed passwords required for systems to authenticate users should be readable only with super-user privileges.

7. *Quick wins:* Utilize access control lists 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.
8. *Quick wins:* Through policy and user awareness, require that administrators establish unique, different passwords for their administrator and nonadministrator 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.
9. *Quick wins:* Configure operating systems so that passwords cannot be re-used within a certain timeframe, such as six months.
10. *Visibility/Attribution:* Implement focused auditing on the use of administrative privileged functions and monitor for anomalous behavior (e.g., system reconfigurations during the night shift).
11. *Visibility/Attribution:* Configure systems to issue a log entry and alert when an account is added to or removed from a domain administrators group.
12. *Configuration/Hygiene:* All administrative access, including domain administrative access, should use two-factor authentication.
13. *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 nonadministrative account. Then, once logged in to the machine without administrative privileges, the administrator should transition to administrative privileges using tools such as Sudo on Linux/UNIX, RunAs on Windows, and other similar facilities for other types of systems. Users would use their own administrator accounts and enter a password each time that is different that their user account.
14. *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.

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 super-user 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 should 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.

To enforce the requirement for strong passwords, 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 super-user group, such as a domain administrator. Every 24 hours after that point, the system must alert or send e-mail about the status of administrative privileges until the unauthorized change has been corrected or authorized through a change management process. While the 24-hour timeframes represent the current metric to help organizations improve their state of security, in the future organizations should strive for even more rapid alerting.

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 super-user group (such as a domain administrator group) to verify that an alert or e-mail is generated within 24 hours. After this test, the account must be removed from the group and disabled.

Finally, on a periodic basis, the evaluation team must run a script that determines which browser and e-mail client programs are running on a sample of 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 System Entity Relationship Diagram (ERD)

Organizations will find that by diagramming the entities necessary to fully meet the goals defined in this control, it will be easier to identify how to implement them, test the controls, and identify where potential failures in the system might occur.

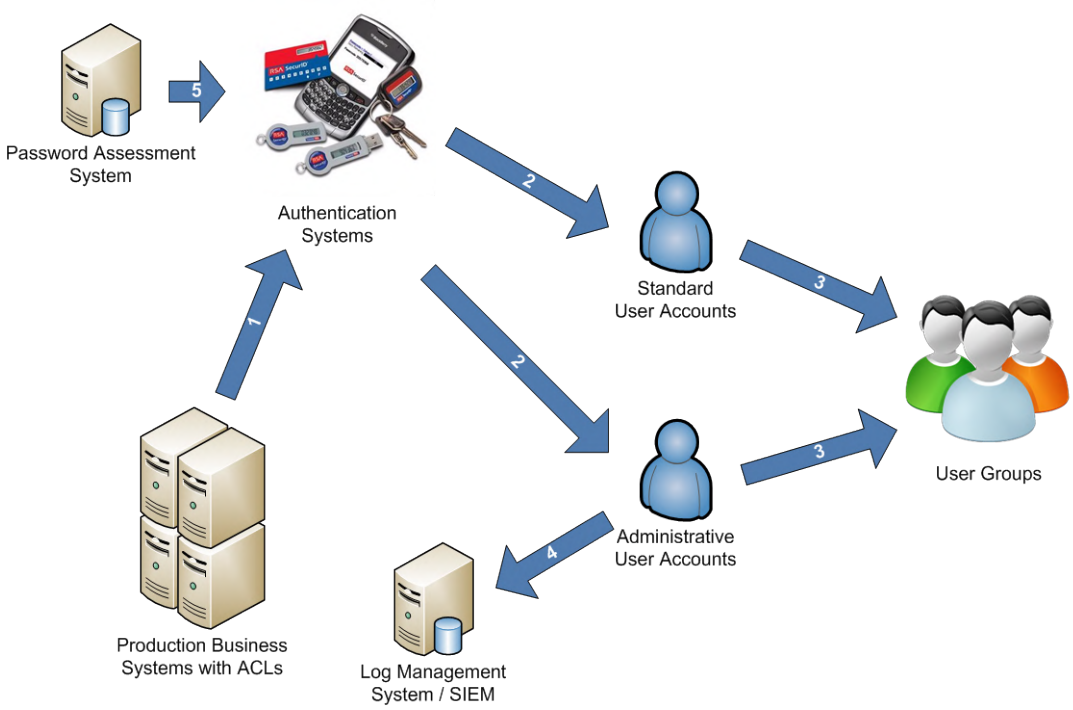

A control system is a device or set of devices used to manage, command, direct, or regulate the behavior of other devices or systems. In this case, we are examining the components of user account provisioning and user authentication. The following list of the steps in the above diagram shows how the entities work together to meet the business goal defined in this control. The list also helps identify what each of the process steps is in order to help identify potential failure points in the overall control.

Step 1: Production systems use proper authentication systems

Step 2: Standard and administrative user accounts use proper authentication systems

Step 3: Standard and administrative user accounts properly managed via group memberships

Step 4: Administrative access to systems properly logged via log management systems

Step 5: Password assessment system validates the strength of the authentication systems

Critical Control 13: Boundary Defense

The processes and tools used to detect/prevent/correct the flow of information transferring networks of different trust levels with a focus on security-damaging data.

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. And despite 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

1. *Quick wins:* Deny communications with (or limit data flow to) known malicious IP addresses (black lists), or limit access only to trusted sites (white lists). Tests can be periodically carried out by sending packets from bogon source IP addresses (unroutable or otherwise unused IP addresses) into the network to verify that they are not transmitted through network perimeters. Lists of bogon 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:* 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) or log analytics system so that events can be correlated from all devices on the network.

3. *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.
4. *Visibility/Attribution:* 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.
5. *Visibility/Attribution:* 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.
6. *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.
7. *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.
8. *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.
9. *Configuration/Hygiene:* 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.
10. *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.
11. *Configuration/Hygiene:* Develop plans to rapidly deploy filters on internal networks to help stop the spread of malware or an intruder.
12. *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
13. *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.
14. *Advanced:* Deploy NetFlow collection and analysis to DMZ network flows to detect anomalous activity.

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

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.

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.

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 System Entity Relationship Diagram (ERD)

Organizations will find that by diagramming the entities necessary to fully meet the goals defined in this control, it will be easier to identify how to implement them, test the controls, and identify where potential failures in the system might occur.

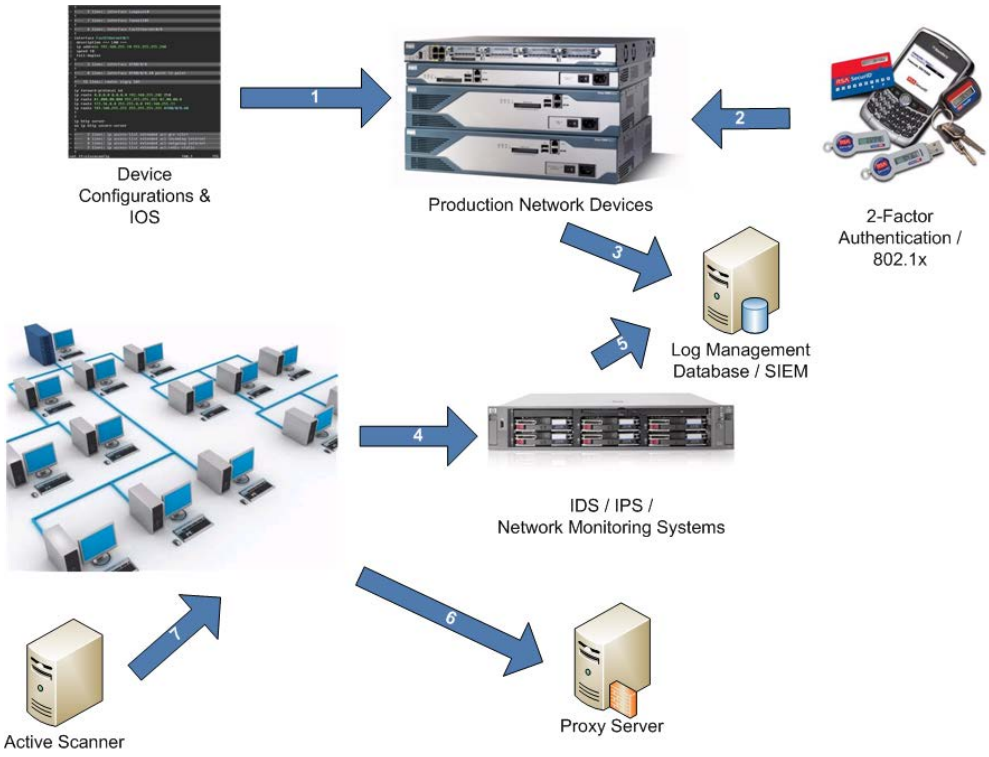

A control system is a device or set of devices used to manage, command, direct, or regulate the behavior of other devices or systems. In this case we are examining the network boundary devices and the supporting systems such as authentication servers, two-factor authentication systems, network monitoring systems, and network proxy devices. The following list of the steps in the above diagram shows how the entities work together to meet the business goal defined in this control. This also helps identify what each of the steps is in order to help identify potential failure points in the overall control.

Step 1: Hardened device configurations applied to production devices

Step 2: Two-factor authentication systems required for administrative access to production devices

Step 3: Production network devices send events to log management and correlation system

Step 4: Network monitoring system analyzes network traffic

Step 5: Network monitoring system sends events to log management and correlation system

Step 6: Outbound traffic passes through and is examined by network proxy devices

Step 7: Network systems scanned for potential weaknesses.

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

The processes and tools used to detect/prevent/correct the use of systems and information based on audit logs of events that are considered significant or could impact the security of an organization.

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:* 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.
2. *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 such a format.
3. *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.
4. *Quick wins:* Develop a log retention policy to make sure that the logs are kept for a sufficient period of time. As APT (advanced persistent threat) continues to stealthily break into systems, organizations are often compromised for several months without detection. The logs must be kept for a longer period of time than it takes an organization to detect an attack so they can accurately determine what occurred.
5. *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.
6. *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.

7. *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.
8. *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.
9. *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.
10. *Visibility/Attribution:* Deploy a SIM/SEM (security incident management/security event management) or log analytic tools for log aggregation and consolidation from multiple machines and for log correlation and analysis. Using the SIM/SEM 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.
11. *Advanced:* Carefully monitor for service creation events. On Windows systems, many attackers use psexec functionality to spread from system to system. Creation of a service is an unusual event and should be monitored closely.

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 SIM/SEM solutions for reviewing logs can provide value, but the capabilities employed to analyze audit logs are quite extensive, including, importantly, even 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.

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 System Entity Relationship Diagram (ERD)

Organizations will find that by diagramming the entities necessary to fully meet the goals defined in this control, it will be easier to identify how to implement them, test the controls, and identify where potential failures in the system might occur.

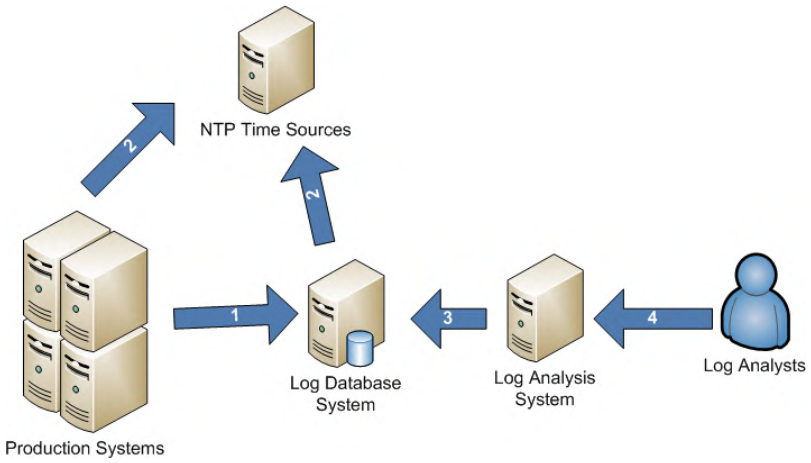

A control system is a device or set of devices used to manage, command, direct, or regulate the behavior of other devices or systems. In this case, we are examining audit logs, the central log database system, the central time system, and log analysts. The following list of the steps in the above diagram shows how the entities work together to meet the business goal defined in this control. It also helps identify what each of the steps is in order to help identify potential failure points in the overall control.

Step 1: Production systems generate logs and send them to a centrally managed log database system

Step 2: Production systems and log database system pulls synchronize time with central time management systems

Step 3: Logs analyzed by a log analysis system

Step 4: Log analysts analyze data generated by log analysis system

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

The processes and tools used to track/control/prevent/correct secure access to information according to the formal determination of which persons, computers, and applications have a need and right to access information based on an approved classification.

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:* Any sensitive information should be located on separated VLANS with proper firewall filtering. All communication of sensitive information over less-trusted networks needs to be encrypted.
2. *Visibility/Attribution:* 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).
3. *Visibility/Attribution:* 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. *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.

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

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 were 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 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.

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. Upon completion of the test, these accounts must be removed.

Control 15 System Entity Relationship Diagram (ERD)

Organizations will find that by diagramming the entities necessary to fully meet the goals defined in this control, it will be easier to identify how to implement them, test the controls, and identify where potential failures in the system might occur.

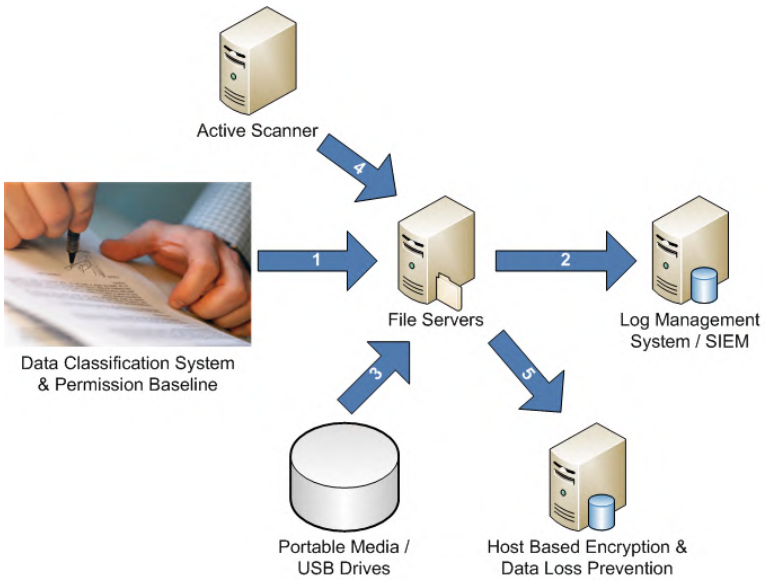

A control system is a device or set of devices used to manage, command, direct, or regulate the behavior of other devices or systems. In this case, the data classification system and permission baseline is the blueprint for how authentication and access of data is controlled. The following list of the steps in the above diagram shows how the entities work together to meet the business goal defined in this control. The list also helps identify what each of the process steps is in order to help identify potential failure points in the overall control.

Step 1: An appropriate data classification system and permissions baseline applied to production data systems

Step 2: Access appropriately logged to a log management system

Step 3: Proper access control applied to portable media/USB drives

Step 4: Active scanner validates, checks access and checks data classification

Step 5: Host based encryption and data loss prevention validates and checks all access requests

Critical Control 16: Account Monitoring and Control

The processes and tools used to track/control/prevent/correct the use of system and application accounts.

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:* All accounts should have an expiration date associated with the account.
3. *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.
4. *Quick wins:* Establish and follow a process for revoking system access by disabling accounts immediately upon termination of an employee or contractor.
5. *Quick wins:* Regularly monitor the use of all accounts, automatically logging off users after a standard period of inactivity.
6. *Quick wins:* Monitor account usage to determine dormant accounts that have not been used for a given period, such as 45 days, notifying the user or user's manager of the dormancy. After a longer period, such as 60 days, the account should be disabled.
7. *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.
8. *Quick wins:* All nonadministrator accounts should be required to have strong passwords that 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. These values can be adjusted based on the specific business needs of the organization.
9. *Quick wins:* Account lockout should be used and configured such that after a set number of failed login attempts the account is locked for a standard period of time.
10. *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.
11. *Visibility/Attribution:* Monitor attempts to access deactivated accounts through audit logging.

12. *Configuration/Hygiene:* 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.

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 System 16 Entity Relationship Diagram (ERD)

Organizations will find that by diagramming the entities necessary to fully meet the goals defined in this control, it will be easier to identify how to implement them, test the controls, and identify where potential failures in the system might occur.

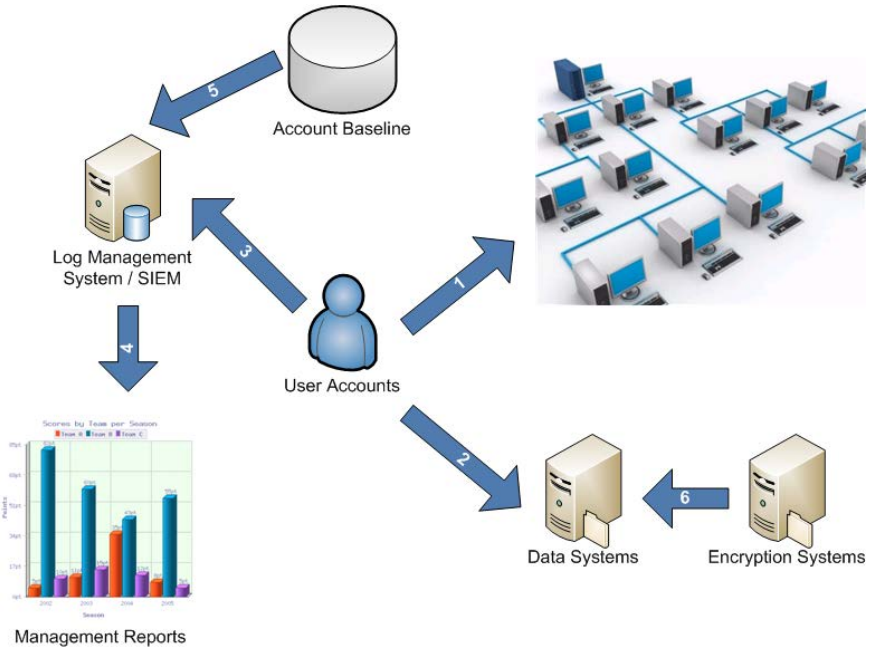

A control system is a device or set of devices used to manage, command, direct, or regulate the behavior of other devices or systems. In this case, we are examining user accounts and how they interact with the data systems and the log management systems. Another key component of these systems is the reports generated for management of user accounts.

The following list of the steps in the above diagram shows how the entities work together to meet the business goal defined in this control. It also helps identify what each of the process steps is in order to help identify potential failure points in the overall control.

Step 1: User accounts are properly managed on production systems

Step 2: User accounts are assigned proper permissions to production data sets

Step 3: User account access is logged to log management system

Step 4: Log management systems generate user account and access reports for management

Step 5: Account baseline information is sent to log management system

Step 6: Critical information is properly protected and encrypted for each user account

Critical Control 17: Data Loss Prevention

The processes and tools used to track/control/prevent/correct data transmission and storage, based on the data's content and associated classification.

How Do Attackers Exploit the Absence of this Control?

In recent years, attackers have exfiltrated significant amounts of often-sensitive data from organizations of all shapes and sizes. 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 the sensitive data was 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. Data loss prevention (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:* Deploy approved hard drive encryption software to mobile devices and systems that hold sensitive data.
2. *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.
3. *Visibility/Attribution:* Conduct periodic scans of server machines using automated tools to determine whether sensitive data (i.e., personally identifiable information, 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.

4. *Configuration/Hygiene:* Data should be moved between networks using secure, authenticated, and encrypted mechanisms.
5. *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.
6. *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.
7. *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 be able to detect rogue connections, terminate the connection, and remediate the infected system.
8. *Advanced:* Block access to known file transfer and e-mail exfiltration websites.

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.

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 System 17 Entity Relationship Diagram (ERD)

Organizations will find that by diagramming the entities necessary to fully meet the goals defined in this control, it will be easier to identify how to implement them, test the controls, and identify where potential failures in the system might occur.

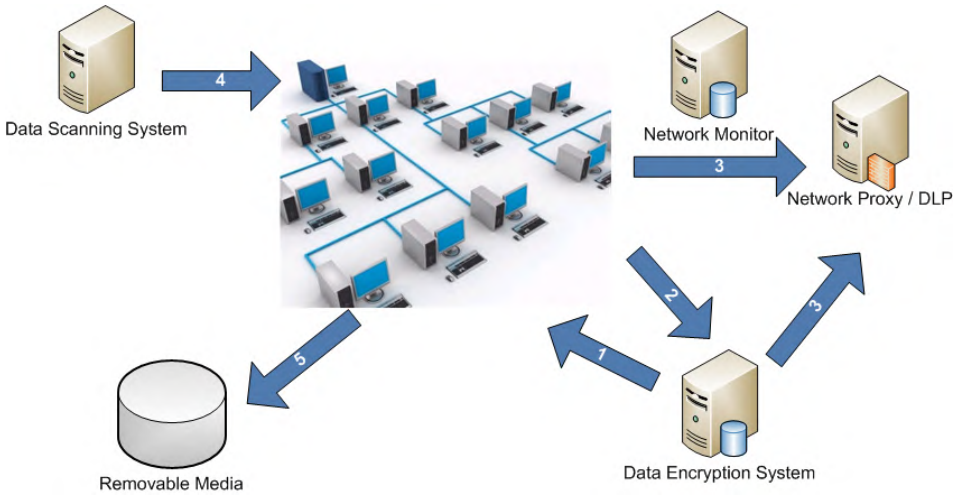

A control system is a device or set of devices used to manage, command, direct, or regulate the behavior of other devices or systems. In this case, we are examining the flow of information in and out of the organization in an attempt to limit potential data loss via network or removable media sources. The following list of the steps in the above diagram shows how the entities work together to meet the business goal defined in this control. It also helps identify what each of the process steps is in order to help identify potential failure points in the overall control.

Step 1: Data encryption system ensures that appropriate hard disks are encrypted

Step 2: Sensitive network traffic encrypted

Step 3: Data connections monitored at the network's perimeter by monitoring systems

Step 4: Stored data scanned to identify where sensitive information is stored

Step 5: Offline media encrypted

Critical Control 18: Incident Response and Management

The process and tools to make sure an organization has a properly tested plan with appropriate trained resources for dealing with any adverse events or threats of adverse events.

Note: This control has one or more sub-controls that must be validated manually.

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.

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

1. *Quick wins:* Ensure that there are written incident response procedures that include a definition of personnel roles for handling incidents. The procedures should define the phases of incident handling.
2. *Quick wins:* Assign job titles and duties for handling computer and network incidents to specific individuals.
3. *Quick wins:* Define management personnel who will support the incident handling process by acting in key decision-making roles.
4. *Quick wins:* 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 Community Emergency Response Team in accordance with all legal or regulatory requirements for involving that organization in computer incidents.
5. *Quick wins:* Assemble and maintain information on third party contact information to be used to report a security incident (i.e., maintain an e-mail address of security@organization.com or have a web page http://organization.com/security).
6. *Quick wins:* 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.
7. *Configuration/Hygiene:* 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 System Entity Relationship Diagram (ERD)

Organizations will find that by diagramming the entities necessary to fully meet the goals defined in this control, it will be easier to identify how to implement them, test the controls, and identify where potential failures in the system might occur.

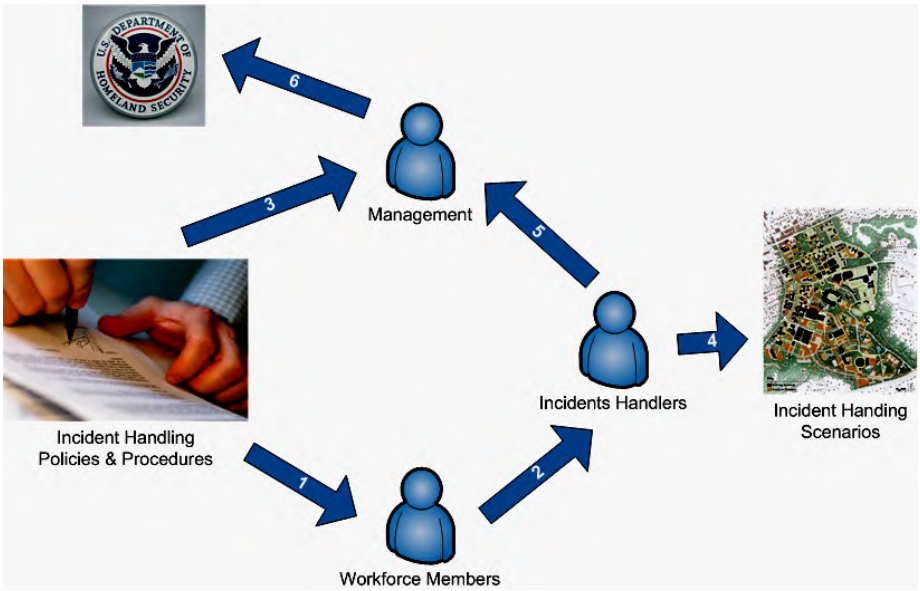

A control system is a device or set of devices used to manage, command, direct, or regulate the behavior of other devices or systems. In this case, we are examining the incident handling process and how prepared organizations are in the event that an incident occurs. The following list of the steps in the above diagram shows how the entities work together to meet the business goal defined in this control. The list also helps identify what each of the process steps is in order to help identify potential failure points in the overall control.

Step 1: Incident handling policies and procedures educate workforce members as to their responsibilities during an incident

Step 2: Some workforce members designated as incident handlers

Step 3: Incident handling policies and procedures educate management as to their responsibilities during an incident

Step 4: Incident handlers participate in incident handling scenario tests

Step 5: Incident handlers report incidents to management

Step 6: Organization's management reports incidents to outside law enforcement and the appropriate computer emergency response team (CERT), if necessary

Critical Control 19: Secure Network Engineering

The process and tools used to build, update, and validate a network infrastructure that can properly withstand attacks from advanced threats.

Note: This control has one or more sub-controls that must be validated manually.

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. *Configuration/Hygiene:* 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 System 19 Entity Relationship Diagram (ERD)

Organizations will find that by diagramming the entities necessary to fully meet the goals defined in this control, it will be easier to identify how to implement them, test the controls, and identify where potential failures in the system might occur.

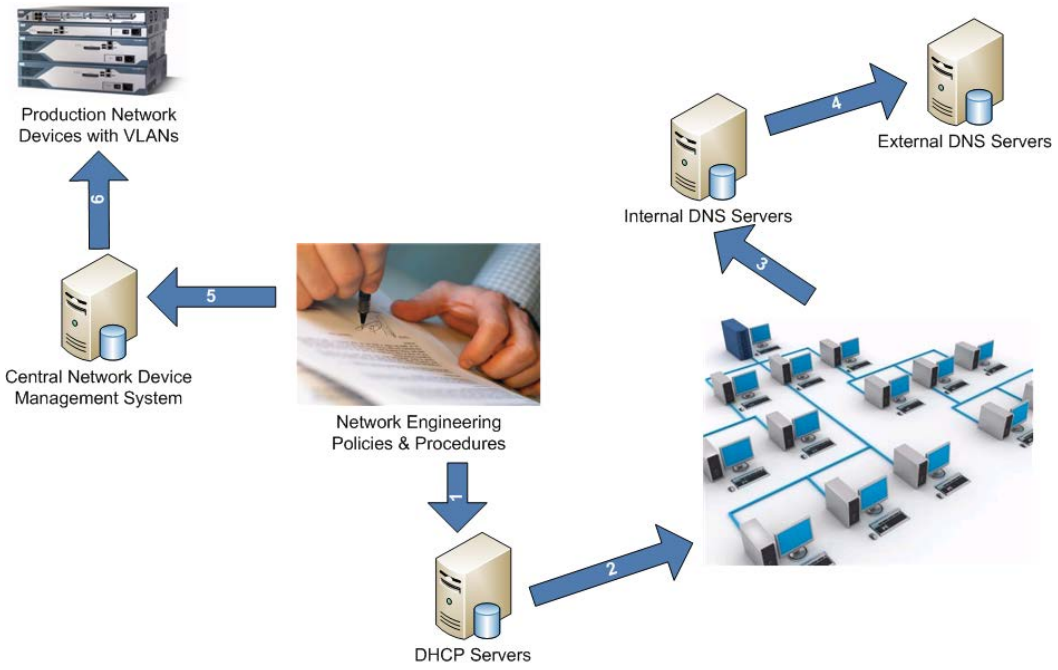

A control system is a device or set of devices used to manage, command, direct, or regulate the behavior of other devices or systems. In this case, we are examining the network engineering process and evaluating the controls that work together in order to create a secure and robust network architecture. The following list of the steps in the above diagram shows how the entities work together to meet the business goal defined in this control. The list also helps identify what each of the process steps is in order to help identify potential failure points in the overall control.

Step 1: Network engineering policies and procedures dictate how network systems function to include DHCP (dynamic host configuration protocol) servers

Step 2: DHCP servers provide IP addresses to systems on the network

Step 3: Network devices perform DNS lookups to internal DNS servers

Step 4: Internal DNS servers perform DNS lookups to external DNS servers

Step 5: Network engineering policies and procedures dictate how a central network management system functions

Step 6: Central network management systems configure network devices

Critical Control 20: Penetration Tests and Red Team Exercises

The process and tools used to simulate attacks against a network to validate the overall security of an organization.

Note: This control has one or more sub-controls that must be validated manually.

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:* 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. *Quick wins:* If any user or system accounts are used to perform penetration testing, those accounts should be carefully controlled and monitored to make sure they are only being used for legitimate purposes.
3. *Visibility/Attribution:* Perform periodic red team exercises to test the readiness of organizations to identify and stop attacks or to respond quickly and effectively.
4. *Visibility/Attribution:* Ensure that systemic problems discovered in penetration tests and red team exercises are fully tracked and mitigated.
5. *Visibility/Attribution:* 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.

6. *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.
7. *Visibility/Attribution:* Plan clear goals of the penetration test itself with blended attacks in mind, identifying the goal machine or target asset. Many APT-style attacks deploy multiple vectors, often social engineering combined with web or network exploitation. Red team manual or automated testing that captures pivoted and multi-vector attacks offers a more realistic assessment of security posture and risk to critical assets.
8. *Configuration/Hygiene:* Use vulnerability scanning and penetration testing tools in concert. The results of vulnerability scanning assessments should be used as a starting point to guide and focus penetration testing efforts.
9. *Advanced:* Devise a scoring method for determining the results of red team exercises so that results can be compared over time.
10. *Advanced:* 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)

Associated NSA Manageable Network Plan Milestones and Network Security Tasks

Milestone 3: Network Architecture

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 System Entity Relationship Diagram (ERD)

Organizations will find that by diagramming the entities necessary to fully meet the goals defined in this control, it will be easier to identify how to implement them, test the controls, and identify where potential failures in the system might occur.

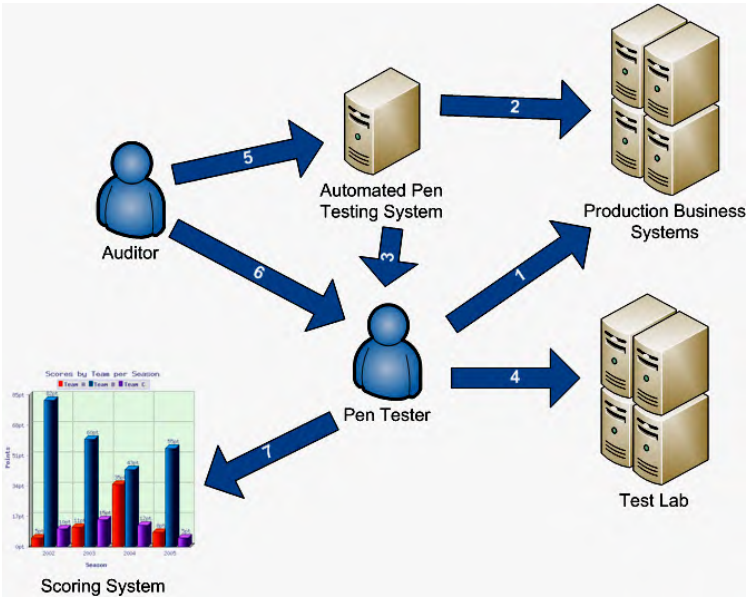

A control system is a device or set of devices used to manage, command, direct, or regulate the behavior of other devices or systems. In this case, we are examining red team and penetration exercises and how those efforts can be valuable to enterprise personnel when identifying which vulnerabilities are present in the organization. The following list of the steps in the above diagram shows how the entities work together to meet the business goal defined in this control. The list also helps identify what each of the process steps is in order to help identify potential failure points in the overall control.

Step 1: Penetration testers perform penetration tests of production systems

Step 2: Automated pen-testing tools perform penetration tests of production systems

Step 3: Automated pen-testing tools inform penetration tester of vulnerabilities discovered

Step 4: Penetration testers perform more extensive penetration tests of test lab systems

Step 5: Auditors evaluate and inspect the work performed by automated pen-testing tools

Step 6: Auditors evaluate and inspect the work performed by penetration testers.

Step 7: Penetration testers generate reports and statistics about the vulnerabilities that have been discovered

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 continue to be updated to reflect what is viewed by the collaborating authors as the most important security controls to defend against cyber attacks.

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 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: 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, 10
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, 5
Attackers compromise target organizations that do not exercise their defenses to determine and continually improve their effectiveness.	4, 5, 11, 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, 15, 17
Attackers scan for remotely accessible services on target systems that are often unneeded for business activities, but provide an avenue of attack and compromise	5, 10, 11

of the organization.	
Attackers exploit weak application software, particularly web applications, through attack vectors such as SQL injection, cross-site scripting, and similar tools.	6, 20
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, 12, 16
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.	9, 12
Attackers exploit boundary systems on Internet-accessible DMZ networks, and then pivot to gain deeper access on internal networks.	13, 19
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 nonsensitive information.	15, 17
Attackers compromise inactive user accounts left behind by temporary workers, contractors, and former employees, including accounts left behind by the attackers themselves who are former employees.	16

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.	12, 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.	15, 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