RAPID7

WHITEPAPER

A Definitive Guide to Understanding and Meeting the CIS Critical Security Controls

The CIS Critical Security Controls are the industry standard for good security. Are you up to par?

Introduction

Everyone in security has heard of the CIS Critical Security Controls, but not all understand exactly how to implement them. The CIS controls supplement almost every other security framework, such as NIST, CSF, NIST 800.53, ISO 27001, PCI, and HIPAA, and they're a useful base for developing or assessing your security program. However, with lengthy documentation and many methods out there for meeting them, implementing these controls can be a daunting project to embark on. Because of this, Rapid7's Advisory Services team, which specializes in security assessments for organizations, developed this guide to explain each control in plain language and assess how it can be approached, evaluated, and implemented.

What are the CIS Critical Security Controls?

The Center for Internet Security (CIS) Top 20 Critical Security Controls (previously known as the SANS Top 20 Critical Security Controls), is an industry-leading way to answer the question on every security practitioner's mind:

"How can I be prepared to stop known attacks?"

Developed by leading security experts from around the world, these controls are not simply a list of best practices to implement, but rather a framework of real ideas and actions gathered from seasoned individuals and enterprises to help everyone strengthen their security posture and stop today's most pervasive and dangerous threats.

The 20 critical controls are divided into three categories: **Basic, Foundational,** and **Organizational**. Basic controls (1–6) should be implemented in every organization for essential defense readiness. Foundational controls (7–16) are the next step up from basic controls, while the Organizational controls (17–20) focus more on people and processes. Whether you're just starting out or are looking to build onto your existing security program, the CIS controls transform best-in-class threat intelligence into prioritized and actionable ways to protect your organization.

The controls were designed in such a way so that they can scale across organizations of any size. Many organizations use the CIS controls as the foundation of their entire security strategy. The sequence of controls allows you to follow a logical path of building your foundation while you gradually improve your security posture and reduce your exposure to risk.

As a subset of the Priority 1 items in the NIST Special Publication 800-53, these controls are also highly relevant and complementary to many established frameworks.
The Rapid7 Advisory Services team relies heavily on the CIS Top 20 Critical Controls as a framework for security program analysis because they are universally applicable to information security and IT governance.

Correct implementation of all 20 of the controls greatly reduces security risk, lowers operational costs, and improves any organization's defensive posture. However, as you likely know, simply being compliant is not enough to entirely mitigate attacks and protect your critical information. While there's no silver bullet for security, organizations can reduce their chances of compromise by moving from a compliance-driven approach to a risk management approach focused on real-world effectiveness.

Implementing the CIS Top 20 Critical Security Controls is a great way to protect your organization from some of the most common attacks.

This guide will help you better understand how to approach and implement each of the key controls so you can go on to develop a best-in-class security program for your organization.

TABLE OF CONTENTS

Control 1

Inventory and Control of Hardware Assets

The theme of this first control is fairly simple: You can't protect what you don't know is there. This means you need the ability to see what is on your network, know which systems belong to whom, and use this information to prevent unauthorized users from connecting to your network. This control is split into eight focused sections relating to network access control, automation, and asset management.

With the ability to actively manage all hardware devices on your network, you can ensure only authorized devices have access and unauthorized ones are found and prevented from gaining access. This requires being able to inventory, track, and correct asset permissions. Implementing inventory control is likely the least glamorous part of a security program, but because it serves as the foundation for many other controls, it can reduce insider-threat and loss risks, clean up the IT environment, and improve the other CIS Critical Security Controls.

How to implement it

Successful implementation often requires bridging existing system inventory or configuration management services with device-based network access control. The inventory management portion is usually based on software or endpoint management services such as the Microsoft System Center Configuration Manager (SCCM), while access control can leverage existing network technology to limit device access to networks.

Implementation of Dynamic Host Configuration Protocol (DHCP) logging and management. will effectively address several sections of Critical Control 1. For organizations with a security information and event management (SIEM) solution or centralized audit repository, ingested DHCP logs can allow correlation with other security and network events. Correlating the logs against additional system information from tools like SCCM or event monitoring services can also assist with inventory tracking and automated inventory management, which leads to added benefits on the financial and operations side of the shop.
If you don't use SCCM, most agent-based system discovery and configuration management can still address this control and other governance requirements.

Though these tools often require time and effort to deploy, the cost benefit is significant, as it allows smaller IT teams to quickly have a major impact on their network.

How Rapid7 can help

The following Rapid7 solutions enhance an organization's ability to discover and identify devices as they connect to corporate assets via the network, email services, or cloud applications:

- InsightVM
- InsightOps
- InsightIDR
- Advisory Services

Rapid7 Global Services will evaluate and document the gaps in your asset discovery process to make recommendations for improving your inventory capabilities.

Control 2

Inventory and Control of Software Assets

Like Control 1, this one addresses the need for awareness of what software is running on your systems and network, as well as finding any unauthorized and unmanaged software and preventing it from being installed or executed. Inventory knowledge and control is an essential security need, and when done correctly, it improves the detection and response aspects of any security program. Because of this, CIS places these controls in the "Top 2" in the same way that the NIST Cybersecurity Framework addresses them as "Priority 1" controls on the 800-53 framework.

How to implement it

To start, Local Administrator access and install rights should not be granted for most users. This limitation also assists with other critical controls that deal with access and authentication.

Once installation rights have been limited, any whitelisting or blacklisting of processes should be done in stages, typically starting with a list of unauthorized applications (a "denylist" or "blacklist") and finishing with a list of authorized applications that comprise the whitelist.
This can be rolled out as an authorized software policy first, then followed up with scanning, removal, and central inventory control. Successful implementations of software inventory control often focus on bridging system configuration management services and software blacklisting and whitelisting. The inventory management portion is usually based on software inventory tools or endpoint management services such as SCCM,
FootPrints, or GPO and local policy controls on Windows.

Beyond limiting administrator and installation rights and blacklisting, you should also set up some form of integrity checking and management. In most cases, this is possible using only OS-based tools, and Microsoft includes integrity management tools in Windows 10. OS-level integrity management tools typically rely on limiting installation based on a list of trusted actors (installers, sources, etc.)
In more comprehensive cases (such as in some endpoint protection services), there are heuristic and behavior-based tools that monitor critical application libraries and paths for change. Because integrity management is intrinsically tied to malware prevention and data protection, implementing this section of the control actually assists with:

- **Control 7:** Email and Web Browser Protections
- **Control 8:** Malware Defenses
- **Control 13:** Data Protection

Getting your inventory in order will also cut down on the amount of work needed when an incident arises and will make policy development and enforcement far easier.

How Rapid7 can help

The following Rapid7 solutions are built to automatically scan systems and catalog the running applications to check them for known vulnerabilities, known malware, and other potential risks:

- InsightVM
- InsightOps
- InsightIDR
- Advisory Services

Control 3

Continuous Vulnerability Management

Understanding and managing vulnerabilities is a continuous activity that requires dedicated time, attention, and resources. Failing to proactively scan for vulnerabilities and address discovered flaws means there is a high likelihood an organization's systems will become compromised.

This control offers guidelines for:

- Performing vulnerability scans
- Monitoring and correlating logs
- Staying on top of new and emerging vulnerabilities and exposures
- Implementing remediation
- Establishing a process to assign risk ratings to vulnerabilities

How to implement it

To begin meeting this control, you need to adopt scanning. CIS states that vulnerability scanning should occur weekly, but that is not always possible due to various circumstances and may depend more so on how mature your organization is from a security standpoint. It is important to have both an internal and external scan — internally facing machines should only have authenticated scans performed on them, and outward-facing devices should have both authenticated and unauthenticated scans performed.

Next, all scanning activities must be logged, monitored, and stored. Your security team must be able to see these events are being generated and then match them to scan logs in order to determine whether the exploit was used against a target known to be vulnerable instead of being part of an actual attack. Scan logs and alerts should be generated and stored to track when and where administrative credentials were being used. This way, you can determine that the credentials are only being used during scans on approved devices and only within approved timeframes.

You then need to understand how vulnerabilities could affect your organization. This control states that there must be a process to risk-rate a vulnerability based on exploitability and potential impact, then use that as guidance for prioritizing remediation. However, what it doesn't spell out is what this process looks like. Here
are three important factors to consider:

1. **Threat level:** What is the importance of the asset in terms of the data it hosts and its exposure level?
2. **Risk of compromise:** What is the likelihood that the vulnerability could compromise this system?
3. **Impact of compromise:** If a particular vulnerability is exploited, how will it affect
the confidentiality, integrity, and availability
of the system and its data?

This rating system can help you determine the order in which to proceed with remediation. To ensure patches are being applied across all systems within the organization, it is recommended to deploy and use an automated patch management tool and software update tool. However, tools are not enough to ensure patches are fully and correctly applied. Vulnerability scans that occur after remediation should be analyzed to check that vulnerabilities that were supposed to be remediated are no longer showing up on the report.

How Rapid7 can help

The following Rapid7 solutions enable continuous data collection from all systems through scanning, integrations, and endpoint agents and simplify remediation workflows in the language of the IT team responsible:

- InsightVM
- InsightAppSec
- Advisory Services

Control 4

Controlled Use of Administrative Privileges

This control can be contentious and is often disliked by system admins and users alike. However, it can have a large impact on risk because it has to do with administrative access.

Two very common attacks rely on privilege to execute. The first is when a privileged user is tricked into opening a malicious email attachment or downloading a malicious file from a website. The second is guessing or cracking an administrator's password to gain access to a target machine. To prevent these common attacks, administrative privilege should be heavily restricted to only those users whose jobs and tasks require it. Regular users should never require admin privileges to conduct daily tasks — even system admins and superusers do not require admin access 100% of the time.

Put another way, much of the effort spent implementing all the other CIS controls can be undone if administrative access is not restricted.

How to implement it

First, you will have to deal with the political issues of doing this. Some users think they need admin access to install software (hint: they don't), some say they need it to do their job, while others simply demand it. While admin rights are required to do some tasks, not all tasks call for them.

Here's an exercise that may be helpful: List the tasks an admin user does on an average day, then mark each task that can be accomplished without admin privileges. Show that list to the person responsible for managing risk in your organization. Then, create a separate, normal user account for admins and require them to use it for all normal tasks. For other tasks, they can escalate into their admin account and then de-escalate when complete. It's an extra step, but it's a secure one.

The conversation

The conversation with your users may be painful, so here are a few pointers. Start by saying, "We're reducing/controlling risk by allowing you to use your privilege only for the tasks that require it." This conveys that privilege is not being taken away from them, but rather needs to be used more securely. For executives who still demand it, point out that they are the highest risk to the organization due to their status.

Technical details

To fully meet the requirements of this control, you'll then need to do the following:

- Change all default passwords on all deployed devices
- Use multifactor authentication for all administrative access
- Use long passwords (14 characters or more)
- Require system admins to have a normal account and a privileged account
- Configure systems to issue alerts on unsuccessful logins to admin accounts
- As an advanced control, require that admin tasks can only be performed on machines that are air-gapped from the rest of the network and only connect to systems they need to administer.

Reducing or controlling admin access is a change to the way things are being done, and fear of change is very powerful. However, by reducing admin privilege and satisfying the first three CIS Critical Security Controls, you can reduce the risks in your organization by 80% or more.

How Rapid7 can help

The following Rapid7 solutions monitor access controls and baseline permitted access to systems in critical environments to identify any suspicious change in settings or behavior:

- Metasploit
- InsightIDR
- Advisory Services

Control 5

Secure Configuration for Hardware and Software on Mobile Devices, Laptops, Workstations, and Servers

This control is about shrinking the attack surface by securing anything that connects to your network.
Default configurations are normally geared toward ease-of-deployment, not security. This includes open and running ports and services, default accounts or passwords, and pre-installed software, which are all exploitable in their default state.

How to implement it

To begin addressing this control, you first have to ask yourself what constitutes a secure configuration. As with most questions in security, the answer is contextual and based on your business rules. Approach this with a mindset of starting as small as you can and gradually opening up your systems and applications until they are usable.

This is great for new systems or those that have yet to be deployed, but what about older systems? You likely can't just shut them down and work through this process. Still, you should seek to reduce the running services and ports, especially those that are known to be vulnerable and not in use.

Here are a few recommended configuration frameworks:

- NIST 800-70 rev 3
- National Vulnerability Database (NVD)
- CIS Benchmarks
- Security Technical Implementation Guide (STIG)

Meeting the requirements of Control 5 also requires addressing these five sub-controls:

1. Document and standardize security configurations for authorized operating systems and software.
2. Maintain secure images or templates for all systems based on approved configuration standards.
3. Store master images and templates on securely configured servers that are monitored 24/7.
4. Deploy system configuration management tools to automatically enforce and redeploy configuration settings to systems on a routine basis.
5. Leverage a Security Content Automation Protocol-compliant configuration monitoring system to verify security configurations, log exceptions, and send alerts when unauthorized changes occur.

The vulnerability management process is continuous, not one and done. Changes to your configurations will occur as systems and applications are patched and updated, new software is introduced, or operational support changes.

How Rapid7 can help

The following Rapid7 products scan existing systems and monitor activity across the modern environment to identify misconfigurations and negative outcomes they may have caused:

- InsightVM
- InsightIDR
- Advisory Services

Control 6

Maintenance, Monitoring, and Analysis of Audit Logs

This control covers a variety of areas, including Network Time Protocol (NTP) configuration, verbose logging of traffic from network devices, best practices for leveraging a SIEM for consolidated view and action points, and how often reports need to be reviewed for anomalies. It also runs alongside or directly connects to many other CIS Critical Security Controls.

How to implement it

Implementation of this control ranges in complexity from a quick win to full configuration. What's your quick win? NTP. By leveraging the various NTP pools that are available (such as those offered by the NTP Pool Project), your systems can check in to a single regionally available server on your network. Because it has obtained its time from the NTP pool, you'll be spared from hours of chasing down information.

Full-scale implementation, on the other hand, requires reviewing log data and issuing alerts on any problems. Most regulatory requirements state that logs should be reviewed "regularly" but remain vague on what this means. A good rule of thumb is to review your logs on a weekly basis at the bare minimum. You'll also have to define what critical alerts look like, who should receive these, and how they should be alerted in order to fine-tune your alerts. SIEM manufacturers and managed service providers have their predefined criteria, but you may have additional needs, so take the time to define your use cases to ensure alerts are sent to the appropriate resources and for the appropriate level of concern to avoid alert fatigue.

Wondering what log data you should be collecting? In a perfect world where storage isn't an issue, each of the following would have security logs sent to the SIEM:

- Network gear: Switches, routers, firewalls, wireless controllers, and their access points
- Third-party security support platforms: Web proxy and filtration, anti-malware solutions, endpoint security platforms (HBSS, EMET), identity management solutions, and IDS/IPS
- Servers: Application, database, web, and file servers, as well as domain controllers in a Windows environment
- Workstations: Pretty much all security log files

This list is by no means exhaustive, so here are a few references to refine what information to include in log collections:

- SANS Log Management Strategies
- NIST SP 800-92
- Malware Archeology

How Rapid7 can help

The following Rapid7 solutions identify systems, ingest audit logs, and identify the anomalies and events of interest for each organization as they occur:

- Metasploit
- InsightIDR
- InsightOps
- Advisory Services

Control 7

Email and Web Browser Protections

The biggest threat surface in any organization is its workstations. This control helps you understand how to manage this threat surface without limiting usability and covers topics such as browser and email client safety, which are critically important for low-level risk mitigation.

How to implement it

Because this control touches on a number of IT functions, it's important to have the people who run the various implicated systems on board when working with it.
Follow these steps:

Start with filtering

Successful implementations usually work from two sides: the server/network side and the endpoint configuration/application side. Networking and email server teams should start by limiting how attachments are handled and forwarded from the mail server to clients, and should implement content filtering first. This is already set up on many mail servers, but it's worthwhile to ensure potentially malicious content is being filtered before it reaches users' inboxes.

Implement SPF, or something similar

Implementing the Sender Policy Framework (SPF) at a DNS level and on mail servers can cut down on the amount of spam and malicious traffic coming into the system. SPF records and implementation should include receiver-side verification. Though a high-effort measure, it's extremely useful for SMTP traffic reduction, better junk-mail sorting, and compatibility with other services.

Configure all the things!

There are a number of ways to handle browser configuration that can both enable your users and limit the risks from malicious code in websites (as well as any attachments that get through your ironclad email server). We typically recommend disabling browser plugins and only running authorized scripting languages and any software that hasn't been reviewed by the security team.

When implementing this control, follow this simple axiom:

You need to make it simple for the users, or they will find a way around it.

Increasing complexity or the effort users have to put in often leads to privilege misuse or other workarounds to defeat the controls.

How Rapid7 can help

The following Rapid7 solutions analyze links received through email, identify filtering services on both endpoints and servers, and test out the effectiveness of all protections:

- InsightVM
- InsightIDR
- Metasploit
- Advisory Services

Control 8

Malware Defenses

This control requires protections at the system, network, and organizational levels. It assesses infrastructure, IoT, mobile devices, and anything else that can be a target for malware — not just endpoints. Control 8 will significantly improve the type of incident response program you're developing. As a bonus, any decent antivirus software still scans for most malware signatures and malicious behavior, and despite claims to the contrary, antivirus is not dead; it just grew up.

How to implement it

Successful implementation of Control 8 requires following these five steps:

1. **Centralize, automate, and configure**
   Centralizing the management of antivirus and anti-malware system logs can greatly accelerate and simplify the incident detection process. You probably already have a good start if you're using any centrally managed antivirus service and managing your workstation and endpoint configuration.

2. **Log your incidents and track them over time**
   Enterprise-level antivirus and anti-malware solutions usually have some form of logging facility, and this — in concert with other logs from firewalls, network instruments, and critical systems — will give the security team a clear picture of what's going on inside the network.
   Log both detection and response information from your antivirus tool, and have a service to monitor the number of infected and damaged machines to help you pinpoint where the malware is.

3. **Anti-malware everything, always**
   Make it clear that everything on your network needs to have an antivirus installed and that anything that is run by your IT team should have an antivirus client that reports back to you. Not only will this give you system-wide visibility, but it will also ensure you are not granting network access to devices that may be carrying malware. Aim for as much coverage as possible.

4. **Enable the detection of OS-level malware, removable media, installation, and tampering**
   Malware can show up anywhere, and removable media is a major source of infection. Your antivirus policy should be able to scan removable media before it's allowed on anything, as well as limit who can install software. Removing root privileges also reduces the risk of user-installed software, malware attacking critical system objects, or exploiting access to admin rights.

5. **Watch your edges**
   Looking at inbound and outbound network traffic from unusual IP addresses will also help with identifying malware patterns as they emerge and informing response activities. IDS and logging play a huge role here — specifically, log session lengths, DNS requests, and traffic patterns to look for access with Command and Control networks used by known malware. Session length logging can also give hints about data exfiltration, and looking at things like failed attempts to authenticate on services may also act as an attack indicator of a virus or worm.

How Rapid7 can help

The following Rapid7 solutions detect both known malware and unknown suspicious software, in addition to testing evasion techniques for malware defenses:

- Metasploit
- Managed Services
- InsightIDR

Control 9

Limitation and Control of Ports, Protocols, and Services

You can greatly reduce your attack surface by knowing what is running on your network and eliminating extraneous means of communication. Control 9 requires that all ports, protocols, and services (PPS) in use within your infrastructure are defined, tracked, and controlled, and that any corrections are made within a reasonable time frame. The initial focus should be on critical assets but should evolve over time to encompass your entire infrastructure.

Most off-the-shelf server software will come with instructions on which ports are required to run the system and allow you to configure things like communications between applications and databases. This can help with creating firewall rules and ensuring your data protection program and disaster recovery/business continuity plans are up-to-date.

How to implement it

Successful implementation of Control 9 requires following these six steps:

1. **Perform a baseline port scan of the hardened system using your vulnerability scanner**
   You can also use other freely available applications, such as port scanners and packet-capturing tools. Once the system is installed, perform another port scan and compare the results. Anything required of the system that wasn't mentioned in the configuration and instructions should be made known at this time.

2. **Leverage host-based firewalls on your servers**
   Whitelists should be configured to only allow communications between aspects of the systems (e.g., database connections or admin access from specific IP spaces). Workstations can use this technology, but other nonessential communications should be blacklisted.

3. **Perform port scans of your infrastructure to understand and control exposure**
   Developing a baseline should be one of the first things you do. When a discrepancy is discovered between the known and approved baseline, your setup should allow stakeholders to receive alerts so they know to investigate the activity and validate its business purpose.

4. **Hit your external IP space by performing port scans against the entire range of external IPs you have assigned** — A number of organizations only scan specific external IP addresses as part of their vulnerability management programs, and there is always a chance that a host may have accidentally been placed in the external space. Finding these hosts and moving them into a VLAN in your internal private IP space is an important part of risk reduction.

5. **Separate critical services on individual host machines** — While you may be leveraging your domain controller for DHCP, you should not include any other critical services on these boxes. Physical segregation is ideal, but in complex computing and operational environments, this may not be feasible. Regardless of segmentation, enhance the security of the hosts by locking them down to only the required services. In the case of critical services such as DNS, DHCP, and database servers, ensure the attack landscape is kept at a minimum and that attackers cannot gain access to multiple lines of advancement to the crown jewels.

6. **Use application firewalls and place them in front of any critical servers** — This helps ensure that only the appropriate traffic is permitted to access the application.

How Rapid7 can help

The following Rapid7 solutions track activity across ports and protocols, identify running services, and test host-based firewalls for susceptibility to attack:

- InsightVM
- InsightIDR
- Metasploit
- Advisory Services

Control 10

Data Recovery Capabilities

Adversaries like to muddle in more than just configurations and software — they can also alter data, potentially contaminating the entire organization. Think about the data in your systems and consider the cascading effects of large-scale contamination, such as the loss of financial reports for a business or health records for a hospital. These types of attacks have the potential to be catastrophic, especially for those handling sensitive information such as PII or medical records.

According to CIS, the following is the key principle of this control:

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

How to implement it

There are several facets to implementation, but policies, processes, and tools related to backups and testing remain central to this control.

Control 10 consists of four criteria:

1. Ensure each system is automatically backed up on at least a weekly basis, and more often for systems storing sensitive information, so that in the event of malware infection, restoration can be from a version that is believed to predate the original infection. All backup policies should be compliant with any regulatory or official requirements.
2. Perform a data restoration test on a regular basis to ensure the backup is working properly.
3. Ensure backups are properly protected via physical security or encryption both at rest and in transit. This includes remote backups and cloud services.
4. Ensure key systems have at least one backup destination that is not continuously addressable through operating system calls. This will mitigate the risk of attacks like ransomware.

Like many things in life, practice makes perfect — or at least reduces the chance of a noteworthy problem or setback. While data recovery procedures make up one small piece of the larger security puzzle, well-rounded data recovery procedures are vital.

How Rapid7 can help

Rapid7’s Advisory Services tailor to your organization’s infrastructure by recommending system backup technology and, most importantly, helping to implement a robust restoration testing process.

Control 11

Secure Configuration for Network Devices Such as Firewalls, Routers, and Switches

The goal of this control is to harden critical network devices against compromise, as well as establish and maintain visibility into both legitimate and malicious changes that occur on them.

How to implement it

Adherence to Control 11 requires the following:

1. **Baselining**
   Compare firewall, router, and switch configurations against standard secure configurations defined for each type of network device. Security configurations should be documented, reviewed, and approved by a change control board, and deviations or updates should be documented and approved.

2. **Change management**
   New configuration rules that allow traffic to flow through network security devices such as firewalls and network-based intrusion prevention systems should be documented and recorded in a configuration management system with a specific reason for each change, who is responsible for the business need, and the duration of the need.

3. **Change detection**
   Use automated tools to verify standard device configurations and detect changes. All alterations to such files should be logged and automatically reported.

4. **Two-factor authentication and encryption**
   Integrate network infrastructure devices with multifactor authentication solutions. Or, consider restricting administration to geographically disparate or independently hosted administrative jump stations and implement two-factor authentication on those stations. This means no telnet, anywhere.

5. **Security updates**
   Install the latest stable version of any security-related updates on all network devices. The days of considering attack surfaces only on your outer boundaries are long gone. A mixture of high-availability configurations for failover, combined with automation for both patching and post-testing, can go a long way in moving beyond this critical security maturity level.

6. **Admin access**
   Network engineers should use a dedicated machine for all administrative tasks or tasks requiring elevated access. This machine should be isolated from the organization's primary network and not allowed internet access. It should never be used for reading email, composing documents, or surfing the web.

7. **Connectivity**
   Manage the network infrastructure across network connections that are separated from the business use of that network, relying on separate VLANs or, preferably, on an entirely different physical connectivity for management sessions for network devices. The easiest and most predictable network to start with is usually your network infrastructure device administration connectivity, where attackers often hide and attempt lateral movement through your environment.

How Rapid7 can help

The following Rapid7 solutions scan existing systems for vulnerabilities and monitor activity across networking devices to identify misconfigurations and suspicious activity:

- InsightVM
- InsightIDR
- Advisory Services

Control 12

Boundary Defense

Today, attackers are constantly probing perimeters for vulnerabilities and information to build their attack plan. However, your boundary defense strategy should not just be about keeping the attackers out, but also keeping sensitive information in. Control 12 defines protections for outside threats, covering your DMZ, firewalls and proxies, IDS/IPS, NetFlow, and remote access.

How to implement it

Use the step-by-step approach below to meet Control 12:

1. **Segment network and control flow**
   Boundary defense is a multilayered approach that requires segmenting your networks in order to control the flow of your data. First, set up a DMZ between your internal network and the internet. Configure it to communicate with the internal network via application layer proxies and set up outbound proxies to filter malicious websites from being visited by end users. Apply blacklists to block traffic to known malicious IPs or whitelists to ban access to everything not needed for approved business purposes.

   To ensure sensitive information isn't being exfiltrated, use proxies to decrypt network traffic and log individual TCP sessions. Set your firewall to block all outbound traffic except for approved business applications. All inbound and outbound traffic should be filtered and monitored.

   Adding to your defense-in-depth strategy, network-based IDS and IPS tools should be implemented. While an IDS issues alerts on attacks by sniffing the traffic flowing through your network, an IPS can actively defend and block unwanted or malicious communications from getting in. Some open source IDS solutions we recommend include Snort, Suricata, and Bro IDS. Most commercial firewall tools offer a network-based IPS.

2. **Collect, analyze, monitor**
   Logs from your firewalls, IPS/IDS, and DMZ should flow through your SIEM solution for correlation, monitoring, and analysis. This centralizes the recording of traffic through the network, network bandwidth, and traffic patterns, streamlines monitoring of applications and protocols that are using the most bandwidth, and detects denial-of-service attacks.

3. **Remote access control**
   When connecting to the internal network, the security profile (configuration policies) should be scanned to ensure security configurations and patch levels are up-to-date. Additionally, all access allowing users to remotely log in to the internal network should require two-factor authentication.

How Rapid7 can help

The following Rapid7 solutions scan perimeter defenses, monitor activity over remote access protocols like VPNs, and effectively test all of these defenses against effective attacker techniques:

- InsightVM
- InsightIDR
- Metasploit
- Advisory Services

Control 13

Data Protection

Data protection is the cornerstone of a solid security program, whether you're following a law to protect certain types of data or regulatory obligations (such as PCI) that require you to make good-faith efforts to keep data safe.

How to implement it

The following are the three types of controls utilized in data protection:

1. **Managerial controls**
   The foundation of a successful implementation begins with executive support for policies that outline what types of data the organization has, how they're classified or organized, and what can and cannot be done with the data. Taking inventory of your data helps you understand your environment and how interconnected systems and subsystems really are. This can also be used to help define data retention requirements and policies.

2. **Procedural controls**
   These controls provide structure and consistency within the organization to protect data. Common procedural controls are performing scans for sensitive information to ensure it is stored where it is supposed to be stored, and developing processes, procedures, and configurations to ensure data is routed and stored in the appropriate areas.

3. **Technical controls**
   These controls actually protect data and include encryption, blocking access to known file transfer and email sites, and blocking USB ports. Data loss prevention (DLP) tools and privileged account management (PAM) tools can also be used to protect data.

   Managerial controls can be the hardest to implement, as they require executive sponsorship, leadership, and funding. Everyone, from the CEO down to the security team, needs to eat the same dog food. Procedural and technical controls are usually easier to put in place, and some can be done for little to no cost, such as blocking USB mass-storage devices or blocking webmail and file transfer websites. And don't forget setting appropriate file and folder permissions and account control lists (ACLs) to restrict access to data to those who have a valid need to know. All of these can provide a great foundational layer of data protection for your organization.

   While policies themselves can't stop a breach or data leakage, they do help employees understand how the organization uses data and what their roles are in protecting that information.

How Rapid7 can help

The following Rapid7 solutions identify passwords and other sensitive data available in plaintext, monitor for exfiltration attempts, and identify the usage of cloud-based file transfer services:

- Metasploit
- InsightIDR
- Advisory Services

Control 14

Controlled Access Based on the Need to Know

This control defines the processes and tools to track, control, prevent, and correct secure access to critical assets such as information, resources, and systems. It's important to establish a formal classification of your data types in order to define which people, computers, and applications have a need and right to access them.

Let's start with some simple yet often unasked questions:

- Do you know which critical assets exist in your organization's network?
- Do you have a data classification policy?
- Who defines the criticality of systems and information?

These are tough questions to answer, but they can be extremely helpful in determining which people, computers, and applications have both the need and right to access critical assets and data.

How to implement it

Use the following step-by-step approach to meet Control 14:

1. Develop an organization-wide data classification policy and apply it to all IT systems. It should include the following levels:
    - **Level 1:** Data for public consumption that may be freely disclosed
    - **Level 2:** Internal data not for public disclosure
    - **Level 3:** Sensitive internal data that could affect the company if disclosed
    - **Level 4:** Highly sensitive corporate, employee, and customer data
2. Segment your network based on the information from your classification policy. For example, all systems with data classified as sensitive should be located on separate VLANs with firewall filtering. AI network switches should enable private VLANs to reduce the ability of an attacker to communicate to other devices on the same subnet from a compromised system.
3. Implement ACLs on all systems and audit not only the ACLs themselves, but also the detailed user access to those systems and data. Follow the principle of least privilege or role-based access control (RBAC) to assign roles to job functions rather than individuals.
4. Encrypt data both at rest and in transit, especially when data traverses trust zones. Your most sensitive data should require secondary authentication in order to access it.
5. Offload and archive old data sets that have not been accessed for a specific length of time.

How Rapid7 can help

The following Rapid7 solutions monitor access controls and baseline permitted access to systems in each customer's environment to identify any suspicious changes in settings or behavior:

- InsightIDR
- Advisory Services

Control 15

Wireless Access Control

Control 15 covers the processes and tools you need to track, control, prevent, and correct the security use of wireless local area networks (WLANs), access points, and wireless client systems. With so many emails, documents, and logins being transmitted, this control has never been more important.

How to implement it

Consider the following step-by-step approach to meet Control 15:

1. **Do not broadcast your SSID**
   While not foolproof, it will stop most curious types from having a peek.

2. **Deploy TLS certificates on your main/secure networks**
   This takes a little extra effort to set up but is far superior because the end-user devices will need the SSL certificate (which you control). This also helps with the threat of rogue access points being set up with the same name.

3. **Use WPA2-Enterprise**
   This forces per-user authentication via RADIUS. Again, it's more involved than setting a shared WPA2 passphrase, but far more secure.

4. **Adjust and limit your radio broadcast levels**
   Some access points are very powerful and may broadcast outside of your building. Tweak these levels to get as close to your building as you can.

5. **Perform wireless (radio frequency) site assessments**
   This can be performed by professional services organizations or by running your own tools to identify rogue wireless devices on your network and verify that the controls you have in place limit wireless to authorized access points.

6. **Create a guest network**
   Having a segmented, bandwidth-limited guest network that does not have access to any critical resources allows your vendors and other visitors to get to their emails and VPNs without you giving them the keys to the kingdom.

7. **Monitor**
   Keeping an eye on (and logging) who is connected to which networks will help in the event of an incident. You know who is in your house, right? This is no different.

   Wireless access is a convenient — and perhaps even mandatory — component of your overall network. The protection of this component should be an elevated and discrete part of any mature security plan. Knowing who is connected and from where is key — it's not like there is a cable to chase down.

How Rapid7 can help

The following Rapid7 solutions identify rogue wireless access points and detect unknown devices connecting to the wireless network to reduce threats from this attack vector:

- Metasploit
- InsightIDR
- Advisory Services

Control 16

Account Monitoring and Control

Control 16 recommends processes to manage the lifecycle (creation, use, dormancy, and deletion) of system and application accounts. You don't need bells and whistles to meet this control, just basic account controls, configuration settings, and two-factor authentication.

How to implement it

Addressing Control 16 requires shoring up your practices around the following three areas:

1. **Account lifecycle management**
   Managing the lifecycle of system and application accounts is one of the most effective controls you can have to protect your organization.
   An in-depth review will help you determine what types of accounts you have, which accounts are still active, and which ones are no longer valid or in use. Work with HR to develop a communication process so that security is made aware when someone is hired, is fired, quits, or goes on a sabbatical. You should ideally be able to disable all access within minutes of an individual leaving the organization. You'll also want to have a policy around how long to keep dormant accounts before they are deleted.

2. **Configuration settings**
   The settings below can have a very positive impact on your security posture:

   1. Automatically log users off after a set of inactivity
   2. Set lock screens on devices
   3. Monitor for stale accounts that may have period fallen through the cracks
   4. Use account lockouts
   5. Set accounts to expire at regular intervals based on business need and risk appetite
   6. Centralize authentication from a single source, such as LDAP

   The first four listed above can be set via Group Policy.

3. **Two-factor authentication**
   Two-factor authentication is one of the most effective controls you can implement to protect your organization, but it has to be done with reasonableness and executive support.
   It should absolutely be used for administrator accounts, dedicated accounts with access to sensitive information, and for remote/VPN access.
   Two-factor authentication is also one of the best bang-for-your-buck controls you can put in place.

How Rapid7 can help

The following Rapid7 solutions audit system authentication controls, test for weak and shared passwords, and alert on any potential authentication-based attacks or misuse of privileges:

- InsightIDR
- Advisory Services

Control 17

Implement a Security Awareness and Training Program

While your users may have a basic understanding of security controls such as antivirus or web filtering controls, they likely don't know the latest defense strategies and what their responsibilities are when it comes to security. This control helps companies implement a program that instructs employees on all of this.

How to implement it

Below is a step-by-step process to meet this control:

Establish baseline security awareness

Create an overview document that outlines the expectations your business has for IT system usage. This could be part of an acceptable use policy, but it doesn't have to be. New hires should be given this document within their first week on the job, and all employees should be reminded of it at least once a year. This can also be done via video trainings to get them up-to-speed fast. The security landscape changes rapidly, so training should be updated regularly so that it's always relevant.

Mature security awareness over time

Next, think about ongoing security awareness training that addresses new technologies, threats, and other requirements employees should be aware of. Consider creating short, quarterly video trainings that touch on the following:

- The latest social engineering tactics (e.g., phishing, phone calls)
- Authentication methods (e.g., two-factor authentication, strong passwords)
- Sensitive data handling
- Mobile, email, and internet security
- Device loss/theft procedures

Put posters like this one around the office, develop phishing awareness training programs, distribute a quarterly security newsletter, and recognize employees who report incidents to reward good behavior.

Zero in on security topics relevant to your business

There are certain security topics applicable to your business that will require additional awareness and training. For example, if you handle healthcare data, there should be training focused on HIPAA and how to handle sensitive information. Training can also be tailored to certain roles, such as support, executives, and sales, or targeted based on best practices your workforce has failed to adhere to. These targeted trainings should be held on a quarterly basis and be mandatory to attend.

Don't patroniz—empower

Don't approach training your workforce as a nuisance — they are your frontline defenders, so it's incredibly important they are trained well and that you develop camaraderie with them. A hostile relationship could upend all your efforts. Let your workforce know that if they accidentally click on a link or lose a device, they can feel safe reporting it to IT. This will empower them to report issues right away. To that end, be sure they know how to reach the IT and ITSEC teams, perhaps by creating a dedicated email account to report issues.

How Rapid7 can help

The following Rapid7 solutions can help you implement a security awareness and training program:

- Phishing Awareness Training
- Advisory Services

Control 18

Application Software Security

This control helps organizations manage the security lifecycle of all software (both developed in-house and acquired) in order to prevent, detect, and correct security issues. Although this control consists of nine sub-controls, many companies (even the ones that are lower on the maturity scale) have at least some components of this control already in place.

Sub-controls include deploying web application firewalls (WAFs), error checking, web application scanning prior to deployment, maintaining separate environments for production and development, hardening of applications, and more.

How to implement it

Control 18 can seem like a lot to cover, but these controls can be met in three steps:

1. **Foster a relationship with application development and procurement groups**
   It's key at this phase to develop a relationship with those working in development, business, change management, and project management. Have a meeting with these stakeholders to get an understanding of how the software development lifecycle (SDLC) works, pain points that can be addressed by security, and regulatory and compliance requirements that are already being addressed. Once you have a good understanding of all of this, you can move on to the next phase.

2. **Use security gates along the way instead of a pass/fail at the end of the SDLC**
   Too many companies wait to do security checks until the end of the SDLC. Instead, security should be addressed along the way using security gates. Work with your development team to identify the security requirements for the products they're responsible for and build a security gate stack to check for them during the SDLC. This way, security can be iterative, not handled at the end.

3. **Ensure you have proper people and tools in place**
   Project owners will be required to make sure implementation and adherence to the requirements in this control. In meeting with stakeholders, ask those who express the most interest to be a security advocate, or ask if they know of someone on their team who is interested. Basic training may be required to get them up to speed — but once that's done, you can bring in tools to ease the burden of implementing your security control around your SDLC, such as a WAF, application firewalls for non-web-facing applications, and even host-based WAFs, as mentioned earlier.

How Rapid7 can help

The following Rapid7 solutions can scan custom applications, third-party software, and databases to identify vulnerabilities and produce clear remediation recommendations:

- InsightVM
- InsightAppSec
- Advisory Services

Control 19

Incident Response and Management

The key principle of Control 19 is to protect the organization's information and reputation in the event of an incident. By developing and implementing an incident response infrastructure, companies can quickly discover new attacks, effectively contain the damage, eradicate the attacker's presence, and restore the network and systems.

How to implement it

Below are four steps to follow when implementing and managing an effective incident response strategy:

1. **Craft plans and procedures**
   The foundation of your incident response strategy is your incident response plan (IRP) and procedures. Create one if you don't have one, but if you do, now is a good time to review and update it. Your IRP should define different procedures (runbooks) for various types of incidents, including malware, data loss, and distributed denial-of-service (DDoS) attacks. Here's a tutorial on how to draft your plan.

2. **Build your team**
   Your IRP should define personnel roles for handling incidents, including the people managing your security and application log information, support personnel, system administrators, and network engineers.
   They may require specific training to ensure they know their role in the investigation when the plan is activated. In addition, bring in legal, public relations, communications, senior leadership, HR, supply management, and vendors.

3. **Test!**
   Testing is critical to ensure your plan will be effective during an actual incident and that employees understand current threats, risks, and their responsibilities in supporting the incident handling team.
   Tabletop exercises and functional/simulation exercises are recommended.

4. **Raise awareness**
   Everyone should be aware of how to report an incident to your organization's security group. Develop standards for how quickly issues should be reported, how to report them, and what information to submit. Then, include this as part of your routine employee awareness activities and make reporting easily accessible by establishing an email address or web page specifically designated for security incidents.

How Rapid7 can help

The following Rapid7 solutions test existing incident response capabilities and ease the detection and response process, optionally through technology or a managed service:

- Metasploit
- InsightIDR
- Incident Response

Control 20

Penetration Tests and Red Team Exercises

Last but certainly not least, Control 20 ensures the overall strength of your defenses by simulating the objectives of an attacker through penetration testing and Red Team exercises.

Penetration testing involves leveraging techniques used by computer attackers to identify vulnerabilities and exploit them. Many organizations fail to perform pen tests out of fear of what will be found, but it's better to know your weaknesses than to discover you were breached through an unpatched vulnerability that went unnoticed.

Red Team exercises are designed for more mature organizations that have been through multiple penetration tests, have remediated vulnerabilities identified during those assessments, and are ready to test their whole organization's security posture through a simulated attack. These engagements simulate a skilled and motivated attacker interested in compromising their specific organization in order to achieve a specific objective, such as gaining access to credit card numbers or sensitive files. Rather than the approach of finding many vulnerabilities and attack paths as possible as in the case of penetration testing, Red Team exercises are meant to test the detection and response capabilities of the client in order to identify gaps in coverage and help direct future investments.

How to implement it

Use the following guidelines to implement effective testing methods:

Define the type of penetration test you need

Not all pen tests are created equal, so it's important to develop clear goals from the outset. Ask yourself these questions before engaging in these activities:

- Are you looking for just a network-based pen test that searches for OS and host-based vulnerabilities or do you want to test your whole organization's preparedness?
- Do you want to evaluate your perimeter defenses (external pen test) or look at your internal defenses based on a presumed compromise (internal network pen test)?
- Do you want to test employee security awareness through social engineering?
- Do you want the pen test team to target a particular section of your network?
- Do you need the team to exclude any systems from their tests?
- Does your organization have many web-based services or applications that could benefit from a web app pen test?
- Do you want to focus on vulnerabilities or do you want to test your detection and response capabilities?

Other considerations to keep in mind

A pen test will often require the use of a system account to perform some authenticated parts of the testing, so you'll want to ensure these accounts are disabled once testing is over — or, at the very least, ensure any activities on those accounts are isolated to the testing windows. It's also critical to use pen test results in conjunction with vulnerability assessment results. Was a pen tester able to exploit a vulnerability that was identified months ago? Is your organization's patch management program effective? There are many ways a pen test can help identify other programmatic deficiencies in your security program.

How Rapid7 can help

The following Rapid7 solutions simplify penetration testing and Red Team operations and track the results over time to help organizations address issues to help prevent future gaps from arising:

- InsightVM
- Metasploit
- Penetration Testing

That's a Wrap

There is a lot to digest in this guide, but the secret is to prioritize which controls are most applicable to your organization and will have the highest impact. Every company has different security requirements, including the regulations they are required to meet, the customer data they're contractually obligated to protect, culture standards, and more.

Because most companies use these 20 controls to build or solidify their security program, it's important to start where you are. Not all of these controls may be immediately applicable to you, and that's fine. Begin with the ones that will meet the most requirements from the law, your customers, your partners, and your own security standards. Over time, it will become apparent which ones should be implemented next, and by referencing this guide, you can get a jump start on understanding and addressing them.

The good news is that addressing many of these controls becomes easier by leveraging a small set of tools, many of which Rapid7 offers. Click here for a complete breakdown of how Rapid7 can help you meet or enhance each control, or reach out to Rapid7 Advisory Services to ask how we can help simplify this process for you. We assist many organizations of different sizes and industries in maturing their security programs, and we'd be happy to help you, too.

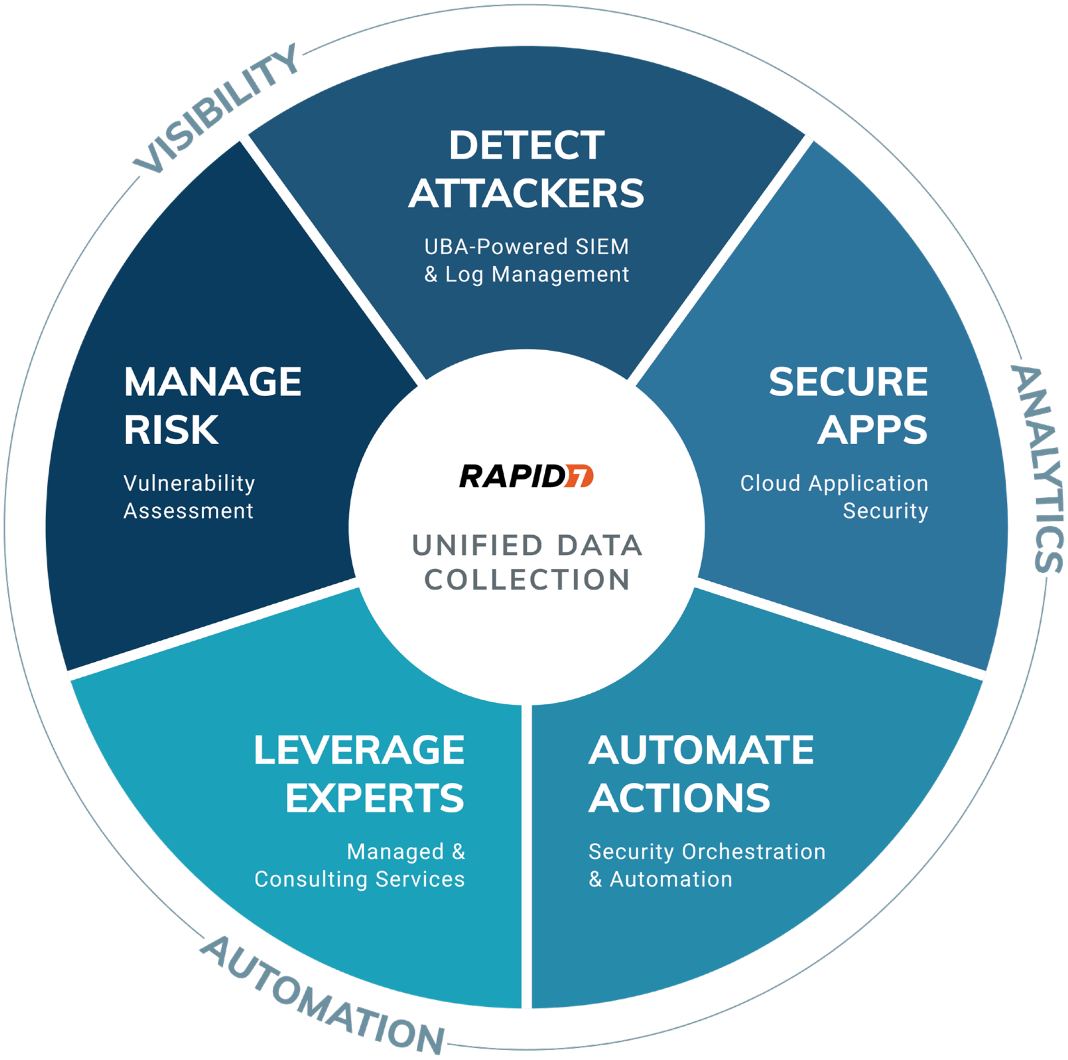

About Rapid7

Rapid7 (Nasdaq: RPD) is advancing security with visibility, analytics, and automation delivered through our Insight cloud. Our solutions simplify the complex, allowing security teams to work more effectively with IT and development to reduce vulnerabilities, monitor for malicious behavior, investigate and shut down attacks, and automate routine tasks. 7,800
customers rely on Rapid7 technology, services, and research to improve security outcomes and securely advance their organizations. For more information, visit our website, check out our blog, or follow us on Twitter.